


OXFORD



LAW FOR
COMPUTER
SCIENTISTS
and OTHER FOLK

MIREILLE
HILDEBRANDT

Law for Computer Scientists and Other Folk

Law for Computer Scientists and Other Folk

Mireille Hildebrandt

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Mireille Hildebrandt 2020

The moral rights of the author have been asserted

First Edition published in 2020

Impression: 1

Some rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, for commercial purposes,
without the prior permission in writing of Oxford University Press, or as expressly
permitted by law, by licence or under terms agreed with the appropriate
reprographics rights organization.



This is an open access publication, available online and distributed under the terms of a
Creative Commons Attribution – Non Commercial – No Derivatives 4.0
International licence (CC BY-NC-ND 4.0), a copy of which is available at
<http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Enquiries concerning reproduction outside the scope of this licence
should be sent to the Rights Department, Oxford University Press, at the address above

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2019952621

ISBN 978-0-19-886088-4 (pbk.)

ISBN 978-0-19-886087-7 (hbk.)

Printed and bound by

CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

for Don Ihde

Acknowledgements

In 2003, Bart Jacobs held his Inaugural Lecture as Professor of Digital Security under the heading ‘Computers under the Rule of Law’. Since then, the institute of Computing and Information Sciences (iCIS) at Radboud University in Nijmegen (the Netherlands) has taken a leading role in integrating legal education in the computer science curriculum. I have had the privilege of teaching law to master students of computer science since 2011, when I was appointed as Professor of ‘Smart Environments, Data Protection and the Rule of Law’.

This has been a very productive experience, and I have found computer science students remarkably open to the foundational grammar and vocabulary of law, and to the architecture of the rule of law. I am deeply grateful for the excellence I encountered, the interesting questions that were raised, and the high quality of the assignments and papers my students submitted. I wrote the draft chapters in the Fall of 2018, during the 2018–19 course (adding some updates during the copy edits at the end of 2019). Two of the students, Aniek den Teuling and Ruben Eijkelenberg, provided salient and extensive feedback.

This book is the consolidation of eight years of teaching law to master students of computer science and I am sure that my new colleague Frederik Zuiderveen Borgesius will find similar intellectual pleasure in teaching. Since I have been awarded an Advanced Grant by the European Research Council for my research project on ‘Counting as a Human Being in the Era of Computational Law’ (COHUBICOL),¹ I am withdrawing from teaching to focus on the research. I am happy to have somehow found the time to turn the course into a book that combines the formats of textbook and scholarly inquiry, very much in line with the core tenets of the COHUBICOL project that has enabled me to publish this work in open access with Oxford University Press (OUP).

COHUBICOL also brought me the pleasure of working with our coordinator Irina Baraliuc, who has performed miracles in the process of getting a draft manuscript online with MIT’s pubpub software for the open review.²

¹ Funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 788734). See <http://www.cohubicol.com>.

² <https://lawforcomputerscientists.pubpub.org>.

viii Acknowledgements

I want to applaud Alex Flach of OUP, who proposed to put this draft online in collaboration with MIT for open review, enabling me to reach out to a more extensive community than the excellent anonymous reviewers that endorsed the book for publication (hoping they will be happy with the final chapter, which deals with some of their comments).

The open review has brought out incisive and detailed comments (online, via email, and over coffee), allowing me to read the text from non-EU perspectives (reminding me of different terminologies, e.g. that of ‘legal power’ in US discourse, instead of ‘legal competence’ in continental European discourse), showing gaps (pointing out that I did not discuss the concept of rights, or that I did not explicitly include the Constitution as a source of law), mistakes (one of my students saliently highlighting an ‘and’ that should be an ‘or’, thus pinpointing a cumulative legal condition that is actually an alternative legal condition), and common misunderstandings amongst non-lawyers who read only one specific chapter (e.g. confusing cybercrime with cybersecurity or private law liability for copyright infringements with cybercrime). Many heart-warming responses also came in via twitter, where the online version has been applauded as a much needed and long-awaited resource, filling a gap that will hopefully contribute to a better understanding between law and computer science. Especially the final chapter on closure, comparing law and ethics, while discussing their relationship and interaction, seems to stir the imagination.

Many thanks are due to my home base in the research group on Law, Science, Technology and Society (LSTS) at Vrije Universiteit Brussels (my main affiliation), with its unique focus on the study of both positive law and legal theory. My long time colleagues Serge Gutwirth and Paul de Hert have steered this research group in myriad ways that foster intellectual independence while acknowledging that research is a *practice* that is always done in collaboration with others, whether in person or through reading (which generates what Montaigne called a ‘monologue interieur’, though a ‘dialogue’ or even a ‘plurilogue interieur’ seem more on-the-spot).

The intellectual shoulders that have carried me this far could fill a universe, if not Borges’ library: lawyers, philosophers of law, philosophers of technology, and computer scientists. To really understand the law, and its foundational relationship with the technologies of written and printed text, one should, however, ‘simply’ read case law. Merely immersing oneself in the intricate reasoning of, for example, the Court of Justice of the European Union or the European Court of Human Rights will generate admiration for those who must think across national jurisdictions, while resolving highly complex

interactions between fundamental rights and freedoms, public interest, private interests, legal certainty, justice, and the instrumentality of the law. I hope this work will contribute to a better appreciation of the role of law in the realm of data- and code-driven environments, based on a proper understanding of the goals, the operations and the limits of the law.

This work is dedicated, however, to Don Ihde, founding father of philosophy of technology in the United States, whose understanding of *how technologies reinvent us as we invent them* is of critical importance to better grasp the assumptions and implications of our new code- and data-driven lifeworld—even if that may not be his primary focus. Don has deeply inspired my appreciation of the systems that computer scientists develop, and the need to inquire how they affect both law and the rule of law.

Mireille Hildebrandt

September 2019
Brussels

Reading Guide

As this is both a textbook and an essay introducing law to computer scientists (and other folk), some guidance on how to read this work seems required.

Please note this book is not an attempt to turn computer scientists into lawyers, there is no claim to completeness. It is a presentation of how *law and the rule of law* protect what is crucial to constitutional democracy and how that is pertinent to computer scientists and other folk. It provides a survey of legal frameworks that apply to developers of computational systems and to those who put such systems on the market or share them otherwise. They should all be aware how their actions may violate the law and what obligations they have. In a constitutional democracy, nobody is above the law.

For *developers* of computational systems, whether based on machine learning, blockchain, or other code, knowledge of the law is also crucial because their systems will co-determine law's effectiveness. If computer systems diminish the substance of human rights or render legal remedies ineffective, they diminish human agency and could even destroy the architecture of constitutional democracy. Before that happens, however, we should expect courts and legislatures to intervene. This is one more reason to pay keen attention to how law operates and to how computational systems can contribute to upholding democracy and the rule of law by protecting the substance of fundamental rights and freedoms.

This work should be read *as a whole* because law is an *architecture* that can only be properly understood if one grasps the whole as well as the parts (including the frictions between them).

However,

- readers not interested in theory *can skip* Chapter 1, and maybe even Chapter 2;
- they will be referred back to the pertinence of these chapters while reading into the parts they deem relevant (in the ebook cross-referencing is supported);
- Chapter 11 is a bonus chapter that targets the intricacies of ethics and code, and how they interact with the law (it can be read together with Chapter 2).

Upfront, please check the glossary, linking the foundational concepts of the law to the sections that use them. Following Wittgenstein, the explanation of these concepts ‘sits’ in the way lawyers *use* them. The glossary thus contributes to a proper understanding of their meaning instead of closing shop by way of (formalizable) definitions.

To preserve the textbook character of the work, footnotes are sparse and only used to refer to relevant sources of law (statutes, case law, treaties), or websites. Each chapter has a concise set of canonical references at the end to enable further reading.

I wish the reader fun, pleasure, and insight. Understanding law is often like solving a puzzle, while simultaneously providing glimpses of how we organize our foundational choice architecture.

Glossary

This glossary orders the **conceptual backbone** of the book, providing the reader with a **vocabulary** and a **grammar of law** as a specific language.

The terms are linked with **the section that introduces or otherwise explains them**. Due to the complexity of the subject no brief definitions are given here. To ‘get’ the meaning the reader will have to ‘mine’ the con-text.

- Absolute rights** 3.1.1
- Agency (legal)** 9.2
- Agency (philosophy)** 9
- Argumentation** 2.1.3
- Automated decisions** 10.3.3.3
- Autonomy**
 - Private law* 3.1.2
 - Privacy* 5.2.1
 - Ethics* 11.1.2
- Blockchain** 10.2
- By design** 10.3
- Competence** 2.1.2.2
- Consent**
 - International law* 4.2.1
 - GDPR* 5.5.2.5; 5.5.2.7; 5.5.2.8
 - Smart contracts* 10.2.2; 10.3.3.3
 - IP law* 7.1
 - Cybercrime* 6.2.1.3
- Constestability**
 - Text-driven law* 1.3; 1.4
 - Legal reasoning* 2.1.3
 - Private law* 3.2
 - Administrative law* 3.3.1.2
 - Criminal law* 3.3.2.2
 - Machine learning* 10.1.3
 - Smart contracts* 10.2.1
 - Automated decisions* 10.3.3.3
 - Legal protection by design* 10.3.2
 - Legal certainty* 11.2.1
 - Fairness by design* 11.3.2
 - Contract law* 3.2.2
 - And smart ‘contracts’* 10.2.2
- Copyright** 7
- Criminal offence** 3.3.2.1
- Criminal procedure** 3.3.2.2
- Cybercrime** 6
- Cybersecurity**
 - And privacy* 5.2
 - CIA* 6.1
 - EU law* 6.3
- Default law** 3.1.2
- Directive (EU)** 4.3.2
- Discrimination** 11.3
- Distributed ledger technologies** 10.2
- Dualism (in international law)** 4.2
- Ethics**
 - Utilitarianism* 11.1.1
 - Deontological ethics* 11.1.2
 - Virtue ethics* 11.1.3
 - Pragmatist ethics* 11.1.4
- Explainability (technical)** 10.1.3
- Fairness**
 - GDPR* 5.5.2.6
 - Ethics* 11.1.2
 - By Design* 11.3
- Formalisation** 11.3.1
- Hart** 2.2.1
- Horizontal effect** 5.3.1
- Human rights law** 5.1
- ICI** 1.3
- Incomputability** 11.1.3
- Instrumentality (of law)** 2.2.2; 11.2
- Intellectual property** 7
- International law** 4
- Interpretation (legal)** 2.1
 - Text-driven law* 1.3; 1.4
 - Explainable ML* 10.1.3
 - In DLTs* 10.2
 - In LPbD* 10.3
 - Jurisdiction* 4.1

xiv Glossary

Justice

As fairness 11.1.2, 11.2.1

As equality 2.2.2

Lawfulness

Administrative law 3.3.1.2

Criminal law 3.3.2.2

GDPR 5.5.2.6; 5.5.2.7

Legal by design 10.3.1

Legal certainty 2.2.2; 11.2.1

In private law 3.1.2

In tort law 3.2.3

Administrative law 3.3.1

Criminal law 3.3.2

Smart contracts 10.2.1

Cp. ethics 11.1.5

By design 11.3.1

'Positivity' of law 2.2.2

Legal conditions 2.1.3

Legal effect 2.1.3

Legal objects 3.1.1

Legal personhood 9

Legal protection

In the onlife world 1.5

ECtHR 5.3.5

CJEU 4.3.3

Legal protection by design 10.3.2

Legal reasoning 2.1.3

Legal remedies

In private and public law 3.1.2

In administrative law 3.3.1.2

For privacy violations 5.4.3

Smart contracts 10.2.1

Legal by design 10.3.1

Legal subjectivity 9.1

Legality principle

In private and public law 3.1.2

In private and criminal law 3.1.3

In constitutional law 3.3.1.1

In administrative law 3.3.1.2

In criminal law 3.3.2

In international law 4.4

Legitimate interest

Legal ground 5.5.2.5

DPIA 10.3.3.1

Liability 8

Machine learning (ML) 10.1

Bias in ML 11.3.2.1

Mandatory law 3.1.2

Micro-targeting 10.1

Monism (international law) 4.2.2

Necessity

Human rights law 5.3.4; 5.3.5

GDPR 4.3.3; 5.5.2.5; 5.5.2.6

Objective law 3.1.3

Obligation

Legal (sources of law) 2.1.1

Legal (nature of legal rules) 2.2.1

Moral (deontology) 11.1.2

Onlife 1.4; 1.6.3; 11.1

Positive law 1; 1.4; 1.5

'Positivity' or 'positiveness' of law 2.2.2

Proportionality test

ECHR 5.3.4

Police access 6.2.2.1

As a balancing test 6.2.2.2

Purpose limitation 5.5.2.6

Radbruch 2.2.2

Regulation 10.2.3

EU Regulation 4.3.2

Techno-regulation 10.3

Relative rights 3.1.1

Representation 9.2

Rights 2.1.2.2

Rules

Primary rules 2.2.1

Secondary rules 2.2.1

Rule of law 2.2

Security

Digital security 5.2; 6.1

Public security 5.3.5

Sensitive data 5.5.2.8

Smart contracts 10.2

Smart regulation 10.2

Sources of law 2.1.1

Sovereignty 4.1.2

Subjective right 3.1.3

Right (see under rights)

Text-driven law 1.4

Tort law 8

Transparency

In constitutional democracy 2.2

In data protection law 5.4.1

GDPR 5.5.2.6

Microtargeting 10.1.3

Smart contracts 10.2.1

Void (smart contracts) 10.2.2

Waldron 4.4; 6.2.2.2

Written law 1.4; 2.1

Table of Contents

List of Abbreviations

xxi

1. Introduction: Textbook and Essay	1
1.1 Middle Ground: Architecture	1
1.2 Law in ‘Speakerspace’	2
1.3 Law in ‘Manuscriptspace’	3
1.4 Law in ‘Bookspace’	5
1.5 Law in Cyberspace: A New ‘Onlife World’	6
1.6 Outline	8
1.6.1 What law does	9
1.6.2 Domains of cyberlaw	9
1.6.3 Frontiers of law in an onlife world	10
1.6.4 Finals	11

PART I WHAT LAW DOES

2. Law, Democracy, and the Rule of Law	17
2.1 What is Law?	17
2.1.1 Sources of law	18
2.1.2 What law does	20
2.1.2.1 Legal effect	20
2.1.2.2 Effective and practical individual rights	25
2.1.3 Legal reasoning	28
2.2 What is Law in a Constitutional Democracy?	31
2.2.1 Law, morality, and politics, and the nature of legal rules	32
2.2.2 Legal certainty, justice, instrumentality	34
3. Domains of Law: Private, Public, and Criminal Law	39
3.1 Private, Public, and Criminal Law: Conceptual Distinctions	39
3.1.1 Absolute rights and relative rights	40
3.1.2 Private law and public law	41
3.1.3 Private law and criminal law	45
3.2 Private Law	47
3.2.1 Property law: transfer of movables	48
3.2.2 Contract law and property law: sale and transfer of real estate	51
3.2.3 Tort liability	54
3.3 Public Law and Criminal Law	57
3.3.1 Public law	58
3.3.1.1 Constitutional law	58
3.3.1.2 Administrative law	58

5.5.2.5	Legal ground for lawful processing of personal data	144
5.5.2.6	Principles of lawful, fair, and transparent processing	148
5.5.2.7	Valid consent	151
5.5.2.8	Special categories of data	153
5.5.2.9	Data protection by design and default (DPbDD)	154
5.5.2.10	Data protection impact assessment	156
5.5.2.11	Compliance and enforcement	158
5.6	Privacy and Data Protection Revisited	160
6.	Cybercrime	163
6.1	The Problem of Cybercrime	164
6.1.1	Computer crime	165
6.1.2	Cybercrime	166
6.2	Cybercrime and Public Law	168
6.2.1	The Cybercrime Convention	168
6.2.1.1	Substantive law	170
6.2.1.2	Procedural law	174
6.2.1.3	Extraterritorial jurisdiction to enforce or investigate	181
6.2.2	Limitations on investigative powers	182
6.2.2.1	Proportionality test for police access to personal data	183
6.2.2.2	Proportionality test, balancing tests, and the image of the scale	184
6.3	The EU Cybercrime and Cybersecurity Directives	187
7.	Copyright in Cyberspace	191
7.1	IP Law as Private Law	192
7.2	Overview of IP Rights	194
7.2.1	Copyright	194
7.2.2	Patents	195
7.2.3	Trademark	197
7.3	History, Objectives, and Scope of Copyright Protection	197
7.4	EU Copyright Law	201
7.4.1	The Copyright Directive and the Enforcement Directive	202
7.4.1.1	The scope of protection (restrictions) and the limitations	202
7.4.1.2	The home copy case of the CJEU	203
7.4.1.3	IP enforcement against intermediaries	204
7.4.1.4	Injunctions to cease unlawful sharing: <i>Sabam v. Netlog</i>	206
7.4.1.5	Injunctions to cease unlawful sharing: <i>Brein v. Ziggo</i>	207
7.4.1.6	The update of the Copyright Directive	210
7.4.2	The Software Copyright Directive	211
7.4.2.1	Exceptions to the exclusionary software copyright: <i>SAS v. WLP</i>	212
7.4.2.2	Exceptions to the exclusionary software copyright: <i>Microsoft</i>	213
7.5	Open Source and Free Access	215
8.	Private Law Liability for Faulty ICT	219
8.1	Back to Basics	220
8.1.1	Chapter 3: private law distinctions	220

xviii Table of Contents

8.1.2	Chapter 4: international and supranational law	222
8.1.3	Chapter 5: data protection law	223
8.2	Tort Law in Europe	225
8.3	Third-Party Liability for Unlawful Processing and Other Cyber Torts	229
8.3.1	Privacy harms	231
8.3.1.1	Canadian ‘tort of intrusion upon seclusion’	231
8.3.1.2	UK ‘tort of misuse of private information’	232
8.3.2	Cyber torts?	233
PART III FRONTIERS OF LAW IN AN ONLIFE WORLD		
9.	Legal Personhood for AI?	237
9.1	Legal Subjectivity	240
9.2	Legal Agency	243
9.3	Artificial Agents	245
9.4	Private Law Liability	246
10.	‘Legal by Design’ or ‘Legal Protection by Design’?	251
10.1	Machine Learning (ML)	252
10.1.1	Exploratory and confirmatory ML research design	253
10.1.2	Implications of micro-targeting	254
10.1.3	Implications of micro-targeting for the rule of law	256
10.2	Distributed Ledger Technologies (DLTs), Smart Contracts, and Smart Regulation	258
10.2.1	Smart contracts and smart regulation	260
10.2.2	The legal status of ‘smart contracts’ under private law	263
10.2.3	The legal status of ‘smart regulation’ under public law	266
10.3	‘Legal by Design’ or ‘Legal Protection by Design’?	267
10.3.1	Legal by design (LbD)	267
10.3.2	Legal protection by design (LPbD)	269
10.3.3	LPbD in the GDPR	270
10.3.3.1	Data protection impact assessment	270
10.3.3.2	Data protection by default and by design (DPbDD)	272
10.3.3.3	Automated decisions	273
PART IV FINALS		
11.	Closure: On Ethics, Code, and Law	283
11.1	Distinctions between Law, Code, and Ethics	284
11.1.1	Utilitarianism and methodological individualism	285
11.1.2	Deontological reasoning: respect for human autonomy	288
11.1.3	Virtue ethics: perceiving the good and doing what is right	291
11.1.4	Pragmatist ethics: taking into account	293
11.1.5	The difference that makes a difference: closure	295
11.2	The Conceptual Relationship between Law, Code, and Ethics	297
11.2.1	Justice, legal certainty, and instrumentality	298
11.2.2	Law, code, and the rule of law	299

11.3	The Interaction between Law, Code, and Ethics	301
11.3.1	'By design' approaches in law and ethics	302
11.3.2	Fairness by design and 'fair computing' paradigms	304
11.3.2.1	The case of COMPAS	306
11.3.2.2	A computational 'fairness by design' approach to detain/release court decisions	310
11.3.2.3	An ethical 'fairness by design' approach to detain/release court decisions	312
11.3.2.4	A legal 'fairness by design' approach to detain/release court decisions	314
11.4	Closure: The Force of Technology and the Force of Law	315

List of Abbreviations

AG	Advocate General
CC	Civil Code
CFREU	Charter of Fundamental Rights of the European Union
CI	contextual integrity
CIA	confidentiality, integrity, availability
CJEU	Court of Justice of the European Union
CoE	Council of Europe
COMPAS	correctional offender management profiling for alternative sanctions
DDOS	distributed denial of service
DLTs	distributed ledger technologies
DPbDD	data protection by design or default
DPD	Data Protection Directive
DPIA	data protection impact assessment
DPO	data protection officer
DRM	digital rights management
EC	European Commission
ECHR	European Convention on Human Rights
ECSC	European Coal and Steel Community
ECtHR	European Court of Human Rights
EDPB	European Data Protection Board
EEC	European Economic Community
EP	European Parliament
ePD	ePrivacy Directive
EU	European Union
FIPs	Fair Information Principles
FSF	Free Software Foundation
FTC	Federal Trade Commission
GCC	German Constitutional Court
GDPR	General Data Protection Regulation (EU)
GPL	General Public Licence
ICCPR	International Covenant on Civil and Political Rights
ICI	information and communication infrastructure
iCIS	institute of Computing and Information Sciences
ICJ	International Court of Justice
ICS	industrial control system
ICT	information and communication technology
iot	internet of things

xxii List of Abbreviations

IP	intellectual property
IPL	international private law
ISPs	internet service providers
LbD	legal by design
LPbD	legal protection by design
ML	machine learning
MLATs	mutual legal assistance treaties
MSs	member states
NCC	Netherlands Criminal Code
NCCP	Netherlands Code of Criminal Procedure
NGOs	non-governmental organizations
OECD	Organisation of Economic Co-operation and Development
OS	operating system
OSI	Open Source Initiative
PCIJ	Permanent Court of International Justice
PDPD	Police Data Protection Directive
PII	personally identifiable information
PSR	Presentence Investigation Report
RIPA	Regulation of Investigatory Powers Act 2000
SARs	subject access requests
TEU	Treaty of the European Union
TFEU	Treaty on the Functioning of the European Union
TRIPs Agreement	Trade-Related Aspects of Intellectual Property Rights Agreement
UN	United Nations
WTO	World Trade Organization

1

Introduction: Textbook and Essay

This book aims to introduce law to computer scientists. For that reason, it serves as a *textbook*, providing an overview of the practice and study of law for a specific audience. Teaching law to computer scientists will always be an attempt, an *essay*, to bridge the disciplinary gaps between two scientific practices that each have their own methodological demands and constraints. This book probes the middle ground, aiming to present a reasonably coherent picture of the vocabulary and grammar of modern positive law (the applicable law in a specified jurisdiction). It is geared to those who have no wish to become lawyers but are nevertheless forced to consider the salience of legal rights and obligations with regard to the construction, maintenance, and protection of computational artefacts. It aims to raise awareness and provide proper information about these legal rights and obligations, not just with regard to computer scientists themselves, but also with regard to those who will suffer or enjoy the results of their constructions. The latter is often considered under the heading of ethics, here it is studied from the perspective of law, explaining the legal rights and obligations involved. It is therefore not a matter of individual moral preferences or intellectual reflection, but a matter of confronting ‘what law does’ when such rights and obligations are violated.

In this introduction I will briefly situate the rise of modern positive law as an affordance of a specific information and communication technology (ICT), namely the printing press, which is even better described as an information and communication infrastructure (ICI). This will be followed by an outline of the book.

1.1 Middle Ground: Architecture

Though many assume that law and computer science are miles apart as scientific disciplines and professional practices, this book takes another position. It is built on the fact that both law and computer science are about architecture, rather than merely about rules (and principles).

Architecture refers to three aspects of both law and computing systems:

1. the fact of being constructed (artificial) rather than natural;
2. the relational and high-dimensional nature of whatever is constructed; and
3. the double ecological nature of the construct
 - as it has to survive in a specific (often dynamic) environment,
 - while the construction itself forms the environment for its inhabitants.

A house, a legal system, and a computing system all have an architecture that determines how the various parts (rooms, legal domains, modules) hold together, interact, and support each other. Architecture refers to physical, institutional, and computational design that determines the strength and sustainability of the construct, involving both hardware (walls, books, silicon chips) and software (the mapping of space to functions such as eating, working, sleeping; the ‘positivity’ or ‘positiveness’ of the law; the program or algorithm). The *high-dimensionality* of the architecture of both law and computer science implies that choices made at any point of the system will ripple through the entire system, resulting in bugs or new features, requiring vigilance as to the dynamics that is inherent in any complex construct, including network effects and unintended consequences. A supreme court that overrules precedent will cause numerous subtle or not so subtle changes in the interpretation of the law by lower courts that need to anticipate how their verdicts will fare. This will in turn trigger adaptations in the conduct of those subject to these courts and may also trigger interventions on the side of legislators or regulators. Law is a complex construct, with a plethora of interlinked, hyperlinked, and deep-linked connections between its various nodes: treaties, statutes, case law, principles, and policies, within and across legal domains such as private law, public law, and criminal law.

1.2 Law in ‘Speakerspace’

Though *we* can hardly imagine what it is like to live in a world without text, the latter is a recent invention. *Homo sapiens* supposedly emerged around 200,000 BC, the script has supposedly been invented around 3,100 BC. Most human societies have thus been oral, meaning that communication was mainly face-to-face. The architecture of ‘speakerspace’ societies is an affordance of human language. The orality obviously limits the reach of language as a means to hold together society, both in space (groupings were

necessarily small) and in time (cross-generational learning depended on word of mouth and durable artefacts). These were non-state, mostly nomadic societies, their livelihood contingent upon hunting (game) and/or gathering (fruits and vegetables).

Anthropologists who spent time in oral societies describe a lifeworld where law, religion, and economy are not merely entangled but non-existent as separable dimensions of society. Clearly, these societies have a *normative order*, they make a difference between interactions that are obligated, preferred, allowed, or prohibited, depending on kinship, age, gender, time of day or year, and context (home, hunting, division of food, celebration, war). This normative order, however, is not externalized in the form of inscriptions on stone, papyrus, or paper. The normativity that rules human interaction in oral society depends on speech and on living memory, aided by a number of mnemonic devices (from rhetorical repetition to artefacts that represent specific taboos or obligations). There is no external written declaration of the norms that govern what is deemed polite, sacrilegious, heroic, expedient, or simply 'proper'. In an oral society, one can neither defend oneself with reference to externalized norms, nor throw them in the face of others. All normativity is, as it were, under the skin of those who are expected to live up to it. This means that the addressants and the addressees of norms are largely the same, requiring repeated assemblies to discuss, establish, and apply such necessarily fluid norms. Being fluid, however, does not imply that such norms are flexible, they may be extremely rigid to compensate for the fluidity of human language (e.g. in the case of taboos) and societal consensus on the existence, interpretation, and application of norms is often delegated to what 'we' (Western anthropologists) like to call priests or others qualified as endowed with special competences.

Note that normativity in oral society mainly depends on the material affordances of the human voice and human memory. There is no police force to implement legal norms and no independent court to contest the way one has been treated; no adjudication apart from negotiated dispute settlement that is based on voluntary jurisdiction.

1.3 Law in 'Manuscriptspace'

As nomadic societies—in the course of centuries—transform into sedentary societies, the relationship with land and time changes due to the need to plough, sow, and harvest. Planning is needed, storage is required, division

of land enacted. The script first emerges as an inscription of numbers, to enable division of land and to count cattle. The rise of the script concerns the rise of an ICI that has far-reaching implications for the size of human society and the way it organizes itself. Sedentary or segmented societies develop into kingdoms and proto-states, with a specialized class of scribes or clerks that holds a factual monopoly on reading and writing. Often, neither the ruler nor those ruled can write or read, and the ruler often governs via his clerks (who are in his service and develop a system of written rules that is used to rule the subjects of the ruler). Note that the role of written 'law' in this era is of two-fold. On the one hand, kings attempt to impose various simple rules (taxes or toll), moving their own position from being a *primus inter pares* (first amongst equals) to being in a position to subject others to their 'general orders backed by threats' (as legal philosopher John Austin famously said). These rules were imposed by a ruler on the ruled, and were e.g. called *capitularia*. On the other hand, kings require their clerks to detect and articulate what is often termed the 'customary law' that rules the relationships between their subjects. The result of this exercise, e.g. the so-called *leges barbarorum*, was used in royal courts as an authoritative though not binding testimony about the applicable law. These rules were not imposed but 'mined' from the oral normativity that supposedly reigned a particular local or kinship group. As with machine learning, the process of 'mining' will inevitably involve framing issues as the norms transition from the management of unwritten expectations to externalized, written records.

The architecture of 'manuscripts space' is an affordance of handwritten manuscripts. The reach of handwritten manuscripts is far beyond that of orality, both in space (the same text can be copied and read across geographical distance) and in time (the text will survive its author and the very same text can be read by later generations). The *distantiation* this involves has curious implications for the interpretation of text; as a text emancipates from the tyranny of its author, its meaning will develop in response to subsequent readers that need to interpret the same text in new circumstances. The rigidity of written manuscripts, so much less ephemeral than spoken words, thus generates a need for iterative interpretation. This also results in the possibility to counter and contest specific interpretations. We can see this 'at work' in the famous medieval version of Roman law, the *Digests*. In the middle of the page, one finds the primary text, as written by Roman juriconsults. On the sides, on the top, and at the bottom, one finds glosses (commentaries) written by medieval lawyers who interpret the primary text in order to apply it to their contemporary society. These glosses were followed, over the course of

centuries, by commentaries on the commentaries, generating a vivid discussion on points of law.

In the end, the stability of *text* combined with the ambiguity of *human language* turns interpretation and contestation into a hallmark of the law, thus offering a very specific type of protection that is at the root of the *legal protection* offered by modern positive law.

1.4 Law in 'Bookspace'

Whereas written manuscripts had to be copied by hand, enabling both error and deliberate changes, the printing press delivered an even more unified text as copies are now 'true' copies. The proliferation of text and the comparative speed of producing identical copies deepen the distantiations in both time and space between text and author, author and reader, and, finally, meaning and text. This intensifies the *quest for stable meaning* in the face of increased opportunities to contest established interpretation. At the same time, the proliferation of printed text (pamphlets, books, newspapers, magazines) invites attempts to systematize content, by way of indexing, developing tables of contents, including footnotes and bibliographies. The architecture of 'bookspace' is more complex, more systematic and hierarchical, and more explicitly interlinked than that of a 'manuscriptspace'. The pressing need for systemization demands taxonomies that are mutually exclusive; books must be categorized in terms of one topic/domain/discipline or another, to enable placing and retrieving them in a private or public library. In his seminal work on information, Gleick explains that *abstract thought is contingent on written text*, as it extends memory and other cognitive resources. Just like the development of counting, calculating, and mathematics depends on notation (for instance, on the invention of 'zero'), abstract thought depends on the sequential processing of written and printed text. This also affords written articulation of more complex frameworks of abstract (general) norms that share the affordances of text-driven abstraction: sequential processing and hierarchical ordering. The combination of the monopoly of violence and the concomitant ability to *impose* abstract legal norms on an abstract population (confined within geographical borders) thus afforded modern positive law: a law explicitly authored by a sovereign that commands obedience from its subjects (internal sovereignty) while protecting them from occupation or interference by other sovereigns (external sovereignty).

This has consequences for the nature of law (which, being artificial, is not fixed):

- sovereigns can now impose general written rules on those subject to their jurisdiction, they can ‘rule by law’;
- sovereigns thereby ‘posit’ the law, which has resulted in ‘positive law’, that is, *a law that is valid in a specific jurisdiction*;
- customary laws are integrated in the legal order of positive law, meaning they must be recognized by the sovereign as valid law (after being ‘mined’ they are imposed);
- the easy proliferation of legal text requires systemization, in the form of elaborate legal codes (in continental law) and treatises (common law) that instigate a complex hierarchy of legal norms, that clarifies which legal norm applies in what situation.

The need for interpretation that is core to text-driven law results in an increasingly independent position for the courts. Originally, judges are appointed by the sovereign to speak the law in his name: *rex est lex animata* (the king is the living law). Kings thus feel free to intervene if a court rules against their wishes. However, as the proliferation of legal text requires study as well as experience, courts increasingly distance themselves from the author of the law (the king), providing a buffer zone between the ruler and those ruled. Montesquieu’s famous *iudex est lex loquens* (the court is the mouth of the law) announces the end of ‘rule by law’ by the sovereign, thus revoking the old adage of *rex est lex animata*. This signifies the beginnings of what we now term ‘the rule of law’, based on an internal division of sovereignty into legislative, administrative, and adjudicative functions that provides for a system of checks and balances. Core to ‘the rule of law’ is an independent judiciary that is capable of sustaining legal certainty, justice, and the instrumentality of the law—if necessary, against the arbitrary will of either the legislature or the administration.

1.5 Law in Cyberspace: A New ‘Onlife World’

One of the challenges that modern, positive law faces, is the transformation of the ICIs of books and mass-media to a digital and computational ICI. Cyberspace refers to *cyber* (steering) and connects with cybernetics (remote control of one’s environment by means of feedback loops). This highlights that the new ICI is fundamentally different from speech, writing, printing, and mass media. Cyberspace is not merely a digitized version of physical space but

refers to an architecture with two novel characteristics: its hyperconnectivity and its computational pre-emptions. In cyberspace the inanimate environment begins to observe, infer, predict, and anticipate human behaviour, while also acting on its own inferences. The ICI does not merely predict the behaviour of its users but also measures and calculates how that behaviour changes when its own behaviour changes (e.g. AB testing). This allows for fine-grained nudging or micro targeting, and for a whole range of automated decisions taken by robotic systems (self-driving cars), the internet of things (domotica), and for governmental and business decisions that directly or indirectly affect individuals or categories of people (behavioural advertising, credit rating, crime mapping, tax fraud detection). The architecture of cyberspace is thus data-driven and code-driven. With the advent of the internet of things (e.g. smart energy grids) and the expected integration of robotics in everyday life (e.g. connected cars) it becomes clear that cyberspace is 'everyware'. Cyberspace is not a separate, virtual space but the emergent architecture of an onlife world. It is *onlife* for two reasons: first, because the difference between online and offline is becoming increasingly artificial, and, second, because the pre-emptive abilities of cyberphysical systems 'animate' our environment. Data-driven infrastructures behave *as if* our environment is alive.

Modern positive law is text-driven. It has developed in an environment driven by text, whose institutional framework is based on text, and whose societal trust and vigilance is contingent on the 'force of law'. Written legal norms are part of a complex legal system that attaches specified legal effect when specified legal conditions apply. Both the conditions and the legal effect are grounded in text and are part of the affordances of human language that are reinforced in printed text. This is related to the fact that speech acts can actually 'do' something, instead of merely describing something. A civil servant who declares a couple 'husband and wife' (or husband and husband, or wife and wife), is not describing a state of affairs but actually 'performs' the marriage. As of that moment the legal effects that private law attributes to a lawful marriage apply, with far-reaching consequences for, for example, inheritance and liability for debts (depending on the applicable national law).

For several centuries, lawyers have been the architects of human societies, structuring economic markets (private law), punitive interventions (criminal law), and the competences of governments to decide crucial matters for their constituents (administrative and constitutional law). In many ways the state itself is a legal construct that defines the contours of everyday life and determines what counts as the public interest. Lawyers may think they still hold a monopoly on the constitution of the state and the foundational structure

of society, but in a society that is increasingly rooted in cyberspace this can no longer be taken for granted. Lawyers now share this ‘monopoly’ with the architects of the internet, the web, and all the different application layers. This especially bears on the *computational backend systems* that are hidden by user-friendly interfaces, while determining the *choice architecture* of their users.

This requires new ways of constructing law. If we value legal protection, we need to articulate it in the data- and code-driven ICI that to a large extent makes and sustains contemporary human societies. This is not an easy quest and it will take some time to achieve anything like it. Time in itself, however, will not do the trick. Just like the rise of the ‘rule of law’ in the era of the ‘bookspace’ was the result of pertinent political struggles, bringing cyberspace under the ‘rule of law’ will require a concerted effort on the side of both lawyers and computer scientists (and, obviously, citizens, policy makers, politicians, and the industry).

In the meantime, it is pivotal that computer scientists get a taste of what law and legal protection is all about, if only to make sure that the systems they study, develop, and maintain are compatible with current legal requirements.

1.6 Outline

As indicated above, computer scientists develop, protect, and maintain computing systems in the broad sense of that term, whether hardware (a smartphone, a driverless car, a smart energy meter, a laptop, or a server) or software (a program, an application programming interface or API, a module, code), or data (captured via cookies, sensors, APIs, or manual input). Computer scientists may be focused on security (e.g. cryptography), on embedded systems (e.g. the internet of things), or on data science (e.g. machine learning). They may be closer to mathematicians or to electrical or electronic engineers, or they may work on the cusp of hardware and software, mathematical proofs, and empirical testing.

Whatever their focus, this book targets ‘law in cyberspace’ from four angles:

1. It answers the question ‘what law is’ by asking the question ‘what law does’.
2. Having introduced the basic elements of the law, this book targets ‘domains of cyberlaw’ that are particularly relevant for computer science: privacy and data protection, cybercrime, copyright, and private law liability.

3. The book discusses the ‘frontiers of law in an onlife world’, notably legal personhood for artificial agents, and legal protection by design.
4. Finally, the closing chapter addresses the relationship between law, code, and ethics, with a focus on algorithmic fairness.

1.6.1 What law does

To prevent mistaking law for either a bag of independent rules or a rigid hierarchical system of decision trees, this book takes off with a discussion of the nature of modern positive law in the light of constitutional democracy, grounding the whole enterprise in a proper understanding of the nature of legal norms and legal reasoning (Chapter 2). This is followed by an introduction of the major legal domains and the logic that informs them (Chapter 3): private law, public law, and criminal law, ending with a basic explanation of international and supranational law (Chapter 4).

These introductory chapters are crucial for a proper understanding of the more targeted legal domains in the second part of the book (on privacy and data protection, cybercrime, copyright, and liability for faulty ICT). The dynamic nature of these targeted legal domains, resulting from the transformative and often volatile nature of our computational lifeworld, requires a foothold in the architecture of modern legal systems.

Without a sound grounding of the core tenets of law and the ‘rule of law’, legal norms are easily subject to misinterpretation and may even contribute to confusion instead of a deeper understanding of how law actually operates.

1.6.2 Domains of cyberlaw

Developing, protecting, or maintaining computing systems will often trigger the applicability of the law, for instance when a software program is protected by copyright or patent, when security breaches are criminal offences, or when default settings are such that data protection law is systematically violated. This provides a practical reason to include law in the *curriculum of computer science* and a good reason to make sure that computer scientists have easy access to concise and correct information about legal domains that are relevant to their work. These legal domains are privacy and data protection (Chapter 5),

cybercrime (Chapter 6), and copyright in cyberspace (Chapter 7), as well as private law liability for faulty ICT (Chapter 8).

This part of the book does not provide a comprehensive in-depth analysis of the domains of cyberlaw. That would take at least four textbooks, if not a proper law degree. The point is not to turn computer scientists into lawyers but to provide them with sufficient information about how these legal domains operate, what kind of questions they should ask when developing computational systems, how to read (often incorrect) headlines on legal issues, and where to find accurate legal information and advice on legal rights and obligations.

1.6.3 Frontiers of law in an onlife world

Next, this book probes three topics on the frontline of law and computer science. First, it investigates the issue of legal personhood for artificial agents (Chapter 9), which refines the understanding of the concept of legal subjectivity and the notion of individual subjective rights. Second, this part of the book examines the concept of legal protection by design (Chapter 10), of which data protection by design is a primary example.

In ‘the old days’—the beginning of this century—an esteemed colleague of mine remarked that my focus on law and computer science was a niche topic for lawyers and legal philosophers. I intuitively guessed that this so-called ‘niche topic’ would come into its own sooner rather than later. Just like international and European law was often considered a niche topic in the 1990s, the relationship between law and computer science will be pivotal for each and every legal domain as each and every practice develops data- and code-driven versions.

By now the tables have turned on lawyers, and they show a growing awareness of the impact of hyperconnected computing systems on the substance of law and on the protections offered by legal procedure. The European Parliament has proposed to consider attributing electronic personhood for certain types of artificial intelligence. The General Data Protection Regulation has imposed a legal obligation to implement data protection by design and default. Law firms, tech start-ups and academia are investing in ‘legal tech’ that some believe will revolutionize the law itself. This book traces the fault lines between modern positive law and its follow-up, arguing that text-driven law offers a type of protection that cannot be taken for granted in an onlife world. The idea, however, is not to reject the new onlife world. The real challenge is to

figure out when to condone it, when to embrace it and when to decline and reject what is on offer.

More precisely, the task is for lawyers and computer scientists to team up and develop a plurality of solutions in close collaboration with those who will suffer and/or enjoy the consequences of the new architecture of our shared world.

1.6.4 Finals

This book ends with a discussion of the distinctions between law, code, and ethics, their interrelationships, and their interaction (Chapter 11). Confusion about the difference between law, regulation, ethics, and policy abounds. Law is not equivalent with regulation, policy is not the same either law or politics. In this volume the issue of closure stands out, because this is what law provides for. Under the ‘rule of law’, however, closure is preceded by potential contestation, and in a democracy closure is performed by a legislature, a public administration, and an independent judiciary acting in concert, based on a set of constitutive checks and balances. All this requires hard work and an acuity as to attempts to achieve closure via other means, either autocratic rule by law or a technocratic rule by technology. In the final chapter I will trace the interactions between code-driven closure, text-driven law, and the space they leave for ethics.

References

Introductions to law at a basic level

Glenn, H. Patrick. 2007. *Legal Traditions of the World*. Oxford: Oxford University Press.
 Hage, Jaap, Antonia Waltermann, and Bram Akkermans, eds. 2017. *Introduction to Law*. 2nd ed. New York: Springer.

Introduction to computer law, information law, information technology law

Bainbridge, David. 2007. *Introduction to Information Technology Law*. 6th ed. Trans-Atlantic Publications, Incorporated.
 Murray, Andrew. 2016. *Information Technology Law: The Law and Society*. 3rd ed. Oxford and New York: Oxford University Press.

On the relationship between law, computers, internet, web, and architecture

- Cohen, Julie E. 2012. *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice*. Yale University Press.
- Hildebrandt, Mireille. 2008. 'A Vision of Ambient Law'. In *Regulating Technologies*, edited by Roger Brownsword and Karen Yeung. Oxford: Hart.
- . 2013. 'The Rule of Law in Cyberspace'. http://works.bepress.com/mireille_hildebrandt/48.
- . 2015. *Smart Technologies and the End(s) of Law. Novel Entanglements of Law and Technology*. Cheltenham: Edward Elgar.
- . 2016. 'Law as Information in the Era of Data-Driven Agency'. *The Modern Law Review* 79 (1): 1–30. <https://doi.org/10.1111/1468-2230.12165>.
- . 2018. 'Law as Computation in the Era of Artificial Legal Intelligence: Speaking Law to the Power of Statistics'. *University of Toronto Law Journal* 68 (1): 12–35. <https://doi.org/10.3138/utlj.2017-0044>.
- Vismann, Cornelia, and Geoffrey Winthrop-Young. 2008. *Files: Law and Media Technology*. Stanford: Stanford University Press. <http://www.loc.gov/catdir/toc/ecip081/2007039414.html>.

On architecture and design in law, politics, and morality

- Lessig, Lawrence. 2006. *Code Version 2.0*. New York: Basic Books.
- Winograd, Terry. 1996. *Bringing Design to Software*. 1st ed. New York and Reading, MA: ACM Press.

On the implications of ICT infrastructures

- Eisenstein, Elisabeth. 2005. *The Printing Revolution in Early Modern Europe*. Cambridge and New York: Cambridge University Press.
- Gleick, James. 2010. *The Information: A History, A Theory, A Flood*. New York: Pantheon.
- Goody, Jack. 1986. *The Logic of Writing and the Organization of Society*. Cambridge and New York: Cambridge University Press.
- Ihde, Don. 1990. *Technology and the Lifeworld: From Garden to Earth. The Indiana Series in the Philosophy of Technology*. Bloomington: Indiana University Press.
- Ong, Walter. 1982. *Orality and Literacy: The Technologizing of the Word*. London and New York: Methuen.

On the move from online and offline to onlife

Floridi, Luciano (ed.). 2014. *The Onlife Manifesto—Being Human in a Hyperconnected Era*. Cham Heidelberg New York Dordrecht London: Springer.

Hildebrandt, Mireille. 2015. *Smart Technologies and the End(s) of Law. Novel Entanglements of Law and Technology*. Cheltenham: Edward Elgar.

On the Rule of Law

Dworkin, Ronald. 1991. *Law's Empire*. Glasgow: Fontana.

Waldron, Jeremy. 2011. 'The Rule of Law and the Importance of Procedure'. *Nomos* 50: 3–31. www.jstor.org/stable/24220105.

PART I

WHAT LAW DOES

This part introduces the conceptual structure of modern positive law, explaining the key concepts of law and the rule of law in terms of what law *does*. This highlights the performative nature of law as a dynamic architecture of legal norms that attribute legal effect whenever specified legal conditions apply. Besides presenting law and the rule of law as a unity of primary and secondary rules that imply a series of foundational legal principles, aiming for a set of antinomian goals, this part also differentiates the main legal domains (public, private, and criminal law) and the concept of jurisdiction as key to the distinction between national, international, and supranational law. In this way, Part I prepares the ground for a better understanding of the vocabulary and the grammar of law that underlies the legal domains that are discussed in Part II (privacy and data protection, cybercrime, copyright, and private law liability).

2

Law, Democracy, and the Rule of Law

Some people believe that law is a set of orders backed by threats, but this raises issues with legal rules that determine when a marriage is valid. Nobody is ordered to marry, there are no threats when you don't. It is just that you cannot get married if you don't follow the rules. The rules do not *regulate* but *constitute* a marriage in the legal sense of that term. Others like to think of law as a bran-tub of social norms, but many social norms are not legal norms. Shaking hands may be a social norm, but in principle it is not required by law. Many assume that law is a system of legal rules, but what does this say for the principles that ground these rules and the policies that refine them? If these principles are not law, and these policies not under the rule of law, what are they? Still others may claim that law is simply a subset of moral rules (those with teeth), but that would imply morality in driving either right or left.

This chapter will squarely face the question of law's 'mode of existence', by asking what law does—and how. This means checking on the sources of law, the nature of legal reasoning, and the question of the relationship between law, democracy, and the Rule of Law.

2.1 What is Law?

'Trying to define law is like trying to hammer a pudding to the wall', wrote legal historian Uwe Wesel. This does not mean that we have no idea what law could be, but rather that our knowledge is implied or tacit knowledge. Such knowledge may lose part of its meaning when translated into the straitjacket of explicit or positive knowledge. In the end, the proof of the pudding is in the eating: 'the life of the law has not been logic but experience' (as Supreme Court justice Oliver Wendell Holmes wrote).

The fluidity of the legal pudding is also the result of the dynamics and complexity of the environment that modern positive law interacts with. This may be a feature, rather than a bug, as the need for iterant interpretation that is core to written law requires flexibility in the face of changing circumstances. *Legal certainty*, one of the core values of the law, is not about fixating the meaning

of legal norms once and for all. Instead, legal certainty targets the delicate balance between stable, *legitimate expectations* and the ability to reconfigure or contest them.

To prevent us from nailing the legal pudding to the wall (a rather unproductive project), legal philosopher Gustav Radbruch defined law in terms of three constitutive values (see below section 2.2.2):

1. legal certainty;
2. justice; and
3. instrumentality.

To qualify as law, a normative framework must aim to sustain, develop, and balance these values—even though they may be incompatible in concrete cases. This requires a combination of analytical thinking, well-developed argumentation, and a keen acuity as to the implications of interpreting the law one way or another. We will return to this point at the end of the chapter.

2.1.1 Sources of law

A source may be a spa that provides refreshing mineral water, an archive to be used for historical research, a witness queried by a journalist, or an encyclopaedia with information about whatever subject or topic. More generally, a source of knowledge refers to where we can find the answer to questions such as: what is the capital of France? where can I find good wine? what is the structure of DNA?

In law, the term ‘source of law’ has a very specific meaning. It refers to both more and less than a source of knowledge *about* the law, as the sources of law are *constitutive of* law. A source of law (1) provides legal norms with authority based on their origin, and (2) makes legal norms binding in their effect. First, it refers to the origin or provenance of *valid legal norms*, that can only be derived from specific sources that thereby give authority to legal norms. For instance, a newspaper article with information about the law is not a source of law, and neither is a Wikipedia article or the website of a law firm. To ensure legal certainty, only a limited set of sources *counts as* sources of law: international treaties, legislation, case law, doctrine, fundamental principles, and customary law. Only these sources provide legal norms with *authority* and

make them *binding* in a specific jurisdiction (either national, international, or supranational).

1. *Treaties* bind the states that have signed and ratified them. They constitute law *between* those states and—depending on the type of treaty—they may also bind citizens and other legal subjects within those states. In Chapter 6 we will look into the binding effects of treaties in more detail.
2. *Legislation* (including a written Constitution) imposes general legal norms on those that share jurisdiction (e.g. within a national state). These norms enact prohibitions and obligations, including obligations not to interfere and rights to such non-interference or rights to specific actions by others. Legislation is binding for all those subject to its jurisdiction. A written *Constitution* has a special status, as it normally defines the powers within the state (legislative, public administration, courts; the relationships between, for example, the national level and sub-national levels, such as a federation and the states, or the central government and provinces and municipalities). Often the Constitution also contains a set of constitutional rights that aim to protect citizens against the state, comparable to human rights and fundamental rights.
3. *Case law* is the result of judgments made by courts. These judgments are simultaneously the result of applying binding legal norms, and a source of legal norms. This is due to the fact that legal norms must be interpreted in the light of the case at hand, which may differ from prior cases—requiring a new interpretation of existing law.
4. *Doctrine* is a body of texts published by lawyers of standing. These texts, restatements, treatises, scholarly articles, or monographs, develop a specific interpretation of a part of the legal framework. This is done either to provide a systematic introduction to and overview of relevant legislation and case law, or to develop a new line of argument with regard to specific issues (e.g. breach of contract in the case of e-commerce, presumption of innocence with regard to predictive policing, consent in data protection law).
5. *Fundamental principles of law* are the principles that are implied in other legal sources, as they inform the applicability and the application of legal norms. They do not function as ‘rules’ that either apply or do not apply, but as an implied philosophy of law that must be taken seriously when deciding the law. One can think of the principle that equal cases must be treated equally and unequal cases unequally to the extent that they are unequal. Or of the principle of fair play in administrative law, meaning

that government agencies should be impartial when deciding on policy and decision-making.

6. *Customary law* (including an unwritten Constitution) is at stake in the absence of written law, when legal subjects (e.g. states in the realm of international law) have acted in a consistent way thus raising legitimate expectations as to how they consider themselves bound. In principle it requires *usus* (a habit of acting in one way rather than another) and *opinio necessitatis* (a shared opinion that this habit is actually based on a duty to act in such a way). Some states do not have a written Constitution, though the powers of the state are nevertheless defined and restricted (as with written Constitutions). In that case the *Constitution* is part of unwritten customary law, and similar to written Constitutions, it has a special status that ranks its binding force above that of other sources of law.

2.1.2 What law does

2.1.2.1 Legal effect

If sources of law are not merely containers of information ‘about’ the law, what are they? What does it mean to say that law actually ‘does’ things? Let us return to a civil servant declaring a man and a woman ‘husband and wife’ and add examples such as a court sentencing a defendant to five years of imprisonment, or a legislature enacting a speed limit. In all these cases ‘the law’ attributes ‘legal effect’ based on specific conditions being fulfilled. When the law speaks (by mouth of the administration, the court, or the legislature) it actually *performs what it says*.

This is a prime example of *speech act theory*, which discriminates ‘locutionary speech acts’ (Tim and Paula are married) from ‘illocutionary speech acts’ (I declare Tim and Paula to be lawfully married). A locutionary speech act is propositional or descriptive (a is p), whereas an illocutionary speech act is performative since it achieves what it declares (I declare a to be p). The ‘achievement’ that is ‘performed’ actually consists of what lawyers call ‘legal effect’: if all legal conditions for a valid marriage are fulfilled, including the declaration by the civil servant or the registration of the marriage in the civil registry, then the legal effects that positive law attributes to a marriage apply. The precise legal effects will depend on national law. For instance, whether or not a marriage entails a community of property by default differs depending on national law. Dutch law—until 2018—had ‘community of property’ as a default, whereas across the legal systems of the United Kingdom there is a

‘separate property system’ by default. In both cases one can sign a prenuptial agreement with a notary public to change the default. In the case of a community of property, the legal effect consists in both partners being liable for debts incurred by the other, meaning their assets can be seized to compensate for debts of their spouse.

The difference between moral norms and habits on the one hand and legal norms on the other, resides in the specificity of legal effect that is not inherent in moral norms or habits. *Legal effect* is not contingent upon the moral inclinations of the person addressed but takes effect depending on the stipulations of *positive law*. In that sense, law is not ‘soft’, and the study of law is not a ‘soft science’. The law has real effects that make a difference in the real world. If murder is defined one way, you may go free, if defined slightly differently, you may go to jail for ten years. In law, definitions have legal effect, they *make a difference that makes a difference*. The reach of such definitions is determined by whoever gets to define the meaning of a norm. Under the Rule of Law, the legislature determines the law but the court has the final say on the meaning of the law. This does not imply that definitions are easy.

In 2012, the US Supreme Court decided the case of *US v. Jones*.¹ The case was about the lawfulness of GPS tracking of a car by the police, after the warrant expired. The question was whether the evidence gathered, thanks to this tracking, was lawfully obtained or had to be excluded as illegally obtained. The defendant claimed that GPS tracking without a valid warrant violates the Fourth Amendment of the US Constitution:

The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

This amendment basically grants people (1) a right to be secure against unreasonable searches and seizures, (2) meaning that searches and seizures require a specified warrant, which can only be granted in case of (3) probable cause. Up until this decision, it was unclear whether the Bill of Rights prohibits GPS tracking unless a warrant has been given. Clearly, when the Bill of Rights was enacted in 1791, GPS tracking did not exist and no such thing was foreseen by

¹ 10-1259 *US v. Jones* (23 January 2012).

its authors. The Supreme Court had to decide whether GPS tracking nevertheless constitutes a search in the sense of the Fourth Amendment, which is considered unreasonable without a valid warrant.

The Court unanimously voted that GPS tracking was indeed a violation of the Fourth Amendment, with the effect that any evidence obtained based on such tracking could not be used. Three types of Opinion can be written by Supreme Court judges to explain their position with respect to a judgment: (1) the *Opinion of the Court*, explaining the reasoning behind the decision, (2) a *concurring Opinion*, explaining the same decision based on another reasoning, and (3) a *dissenting Opinion*, explaining the reasons for dissenting with the majority about the decision. Since the Court was unanimous in its verdict that GPS tracking constitutes a violation of the Fourth Amendment, there was no dissenting Opinion. However, next to the Opinion of the Court, a concurring Opinion was written, endorsing a different underpinning for the same decision.

The Opinion of the Court, written by Justice Scalia, describes the privacy violation in terms of a physical intrusion upon the property of the defendant (the car), relating this to the tort of trespass. What matters here is the violation of a property right. The concurring Opinion of Justice Sotomayor describes the privacy breach in terms of a violation of the reasonable expectation of privacy, which is directly related to the mobility pattern that can be derived from the location data collected by the GPS tracker. Though both Opinions reach the same conclusion, and thus underlie the same decision, the implications for new cases will be different. Whereas the defendant may not care about the reasoning as long as the evidence is excluded, lawyers will be more interested in the reasoning than in the outcome. To the extent that future cases are similar to the case at hand, the reasons given in the Opinion of the Court will determine their outcome. In fact, a lawyer will also be interested in the argumentation of a dissenting Opinion, because these arguments provide reasons that may be relevant in future case law. This is because the Supreme Court may decide to overrule its own previous line of argument, and follow the argumentation of a dissenter (in previous case law). The reasoning of Justice Scalia has a limited reach for other cases, because it seems to require physical trespass upon the property of another. The reasoning of Justice Sotomayor has a broader scope, as it does not depend on such trespass, and instead considers the far-reaching consequences of mobility profiles for the legitimate expectations of privacy. This reasoning could also uproot the so-called 'third-party doctrine' that has severely restricted the right to privacy in the United States, and will be discussed briefly in section 5.2.1.

The legal effect of this judgment is extensive but nevertheless subtle:

1. The decision clarifies that the police need a warrant to place a GPS tracker under a car. This has far-reaching consequences for the practice of policing and obviously for the protection of the privacy of US citizens.
2. If the reasoning of the concurring Opinion gains traction in subsequent Supreme Court decisions, future cases may offer more effective protection in the onlife world.²

In 2014, the Court of Justice of the European Union (CJEU) decided a case that questioned the validity of the Data Retention Directive 2006/24/EC.³ This Directive aims to harmonize the law of the member states (MSs) of the EU, with regard to the retention of telecom data by telco providers. The goal is to ensure that such data remains available for police investigation of serious crime and terrorism. The Directive only concerns metadata, such as traffic data, time-stamped location data, and identification data; it does not require the retention of the content. The Court notes that such metadata provides detailed knowledge of a person's whereabouts and of their relational network, thus enabling very precise insights into a person's private life. The Court concludes that such retention interferes with the fundamental rights to privacy and data protection, as formulated in the Charter of Fundamental Rights of the European Union (CFREU):

Article 7 Respect for private and family life

Everyone has the right to respect for his or her private and family life, home and communications.

Article 8 Protection of personal data

Everyone has the right to the protection of personal data concerning him or her. Such data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right of access to data which has been collected concerning him or her, and the right to have it rectified. Compliance with these rules shall be subject to control by an independent authority.

² Subsequent case law: *Riley v. California*, June 25, 2014, No. 13-132, 573 US, holding: 'The police generally may not, without a warrant, search digital information on a cellphone seized from an individual who has been arrested.' See <http://www.scotusblog.com/case-files/cases/riley-v-california/>. And, *Carpenter v. United States*, June 22, 2018, No. 16-402, 585 US, holding: 'The government's acquisition of Timothy Carpenter's cell-site records from his wireless carriers was a Fourth Amendment search; the government did not obtain a warrant supported by probable cause before acquiring those records.' See <http://www.scotusblog.com/case-files/cases/carpenter-v-united-states-2/>.

³ CJEU, 8 April 2014, C-293/12 and C-594/12 (Digital Rights Ireland).

Notably, the Court considers that not being informed of such interferences will generate a feeling of being constantly surveilled. The fundamental rights of privacy and data protection, however, are not absolute in the sense of having unlimited application. Often, these rights will have to compete with other fundamental rights (for instance, freedom of expression), or with legitimate private and public interests. This requires a delicate and well-argued *balancing act*, as it results in the limitation of a fundamental right. Article 52 of the CFREU stipulates the scope of lawful limitations:

Article 52 Scope of guaranteed rights

Any limitation on the exercise of the rights and freedoms recognised by this Charter must be provided for by law and respect the essence of those rights and freedoms. Subject to the principle of proportionality, limitations may be made only if they are necessary and genuinely meet objectives of general interest recognised by the Union or the need to protect the rights and freedoms of others.

The Court finds that the interference does not adversely affect the essence of the rights, because it does not concern the content. It also finds that, in itself, retention to make metadata available for law enforcement is an objective of general interest. However, the Court considers the measures as enacted in the Data Retention Directive to be *disproportional*, that is, appropriate, but not sufficiently circumscribed to ensure that the interference is actually limited to what is strictly necessary: the scope of the retention measures is undifferentiated, there are no limitations or exceptions; no objective criteria have been stipulated or required to prevent data from being used for anything but the most serious offences; the retention period does not differentiate between categories of data; and, finally, the Court observes that storage outside the EU is not prohibited (which reminds us that this judgment was decided in the aftermath of the Snowden revelations).

The final verdict of the CJEU declared the Data Retention Directive invalid, due to a violation of Articles 7 and 8 of the CFREU. This had sweeping consequences, because it meant that the national laws of the MSs of the EU that were based on the Directive might therefore be unlawful, if they shared the shortcomings of the Directive.

These examples of case law show the complexity of legal issues, the prominent role of legislation as well as case law, and the crucial importance of interpretation and contestation. They also show the *performative nature* of legal norms as they attribute legal effect and potentially transform the world we share.

Legal norms are often explained in terms of a system of legal rules. In the next section, 2.1.3, we will see what this means for *legal methodology*, that is for ‘legal reasoning’. Legal reasoning is based on the idea that if specified legal conditions apply then a specified legal effect is attributed. This raises the question: what types of legal effect are available?

The most obvious legal effect attributed when specific legal conditions apply is that an action or state of affairs is either *lawful* or *unlawful*. This may generate subsequent legal effects, such as the actor being *punishable* or the actor being *liable* to pay damages. Often, the legal effect concerns the attribution of legal obligations and legal rights. If two parties conclude a contract of sale, one party has a legal obligation to pay the price, the other party has a legal obligation to transfer the property of the good. Conversely, one party has the right to obtain ownership, the other has the right to receive the money. In the next paragraph, we will briefly discuss the concept of individual rights.

2.1.2.2 Effective and practical individual rights

The concept of rights is an essentially contested concept, as are most of the terms that ground the generative nature of societal intercourse. Some folk may use the term in a loose way, geared more to moral claims (I have a right to hack into your system if you don’t keep it properly secured) than to the *performative language of legal rights*.

A legal right is a very special ‘thing’, providing a legal subject with specified powers to act in relation to others, or the liberty to ensure that others will refrain from interfering with the object of their right. Though we may intuitively think we know what rights are, attempts to define or analyse them usually end up in complicated framings that generate more problems than they solve. One such attempt is Hohfeld’s infamous typology, which dissects the language of rights, claiming that ‘things’ like property rights are in point of fact bundles of claims, liberties, powers, and immunities:

1. One can have a *claim right* against another person that they act in a specific way, which correlates with that other person’s *duty* to act that way (I sell you a book against a specified price and have a claim right to you paying the price; you have a duty to transfer the property of the book to me).
2. One can have a *privilege* (liberty) against another person that you have the freedom to act in a specified way, which correlates with that other

person having *no claim* that one does not act that way (if I own a book I am free to dispose of it and no other person can claim that I cannot throw it away).

3. One can have *power* (authority, competence) over another person to act in a specified way, which correlates with that other person having a *liability* to act as specified (if I am an employer I can require my employees to perform specified tasks that are part of the job; they are liable to carry out those tasks).
4. One can have an *immunity* against another person with regard to specified actions, which correlates with that other person *not having the authority* to disallow such actions (if I am an employee I have an immunity against my employer requiring me to engage in improper or unlawful behaviour; the employer lacks the authority to make me do this).

The owner of a house has a claim right against the person who rents the house, a liberty against anybody trespassing, authority over the broker who is engaged to sell the house, and an immunity against a neighbour asking that they grow a specific type of tree in their garden.

All this is very interesting, though I am not sure the scheme solves the problems we face in real life. For instance, what if the neighbour claims that the trees you grow take away all the light in their kitchen? Do you have an immunity against their right that you cut the trees, or would invoking such immunity qualify as ‘abuse of right’? Also, many authors have detected inconsistencies, for instance because Hohfeldian terms often have another meaning in positive law. For instance, in tort law the term liability refers to the fact that a tortfeasor is legally responsible for the damage caused, resulting in a duty to pay damages. In Hohfeld’s framework the term has another meaning, as it correlates with a competence rather than with a claim right. We shall therefore not use Hohfeld’s terminology in this work, other than to create awareness that rights and obligations have different meanings, depending on what legal effect the law stipulates for them.

What Hohfeld nevertheless demonstrates can be summarized as:

1. rights always play out in *relationships between legal subjects*, they are based on:
 - a *claim right* of a legal subject against one or more legal subjects (such as a property right, or the right to performance of a contract, the right to have one’s privacy respected by the government); or

- a *competence* of a legal subject with regard to one or more legal subjects (such as the competence of the owner of a good to dispose of one's property as one wishes, the competence of the legislature to enact legislation);
2. these rights *necessarily correspond with*:
- a *duty* for another legal subject to act in a specified way in relation to the rightholder; or
 - *the lack of a right* for another legal subject in relation to the rightholder.

Perhaps more importantly, Hohfeld pays little attention to:

- the difference between, on the one hand, *rights that can be invoked erga omnes (against all)*, such as a property right (this type of rights is also called a right *ad rem*, or an absolute right), and, on the other hand, *rights that can only be invoked ad personam (against specified others)*, such as one's contracting party (this type of right is also called a relative right) (see section 3.1.1);
- the difference between, on the one hand, *a right that one or more others act in a specified way*, such as the right to be paid compensation, and, on the other hand, *a right that others refrain from interfering* with a specified object (one's property or one's fundamental right); the latter right is often called a *liberty or a liberty right*;
- the difference, on the one hand, between rights of *private parties* (natural persons or legal persons) based on *private law*, and, on the other hand, the competences of *public authorities* to enact rules that everyone should follow, to adjudicate and to decide requests based on *public law* (legislature, courts, public administration, police, regulations).

It is crucial to take note of the fact that individual rights that can be enforced against others are a *recent invention* (attributed to Hugo Grotius in the sixteenth century), not a natural attribute of either human beings or human society. For such legal rights to be 'practical and effective' a system of institutions must be developed and sustained that ensures that such rights are upheld against the law of the jungle and against the survival of the fittest. To safeguard rights against arbitrary power we need rules, and to protect rules against arbitrary power we need a rule *of law* instead of a rule *by* (means of) law.

Competences are legal powers that enable a legal subject to lawfully act in a way that impacts the legal status of others.⁴ For instance, the owner of a

⁴ In the United States and the United Kingdom lawyers will speak of legal powers, in continental Europe they speak of legal competences.

good has the legal power to transfer the property; a legislature has the competence to enact binding legislation; a court has the competence to authoritatively decide cases, public administration may have competence to take decisions on building permits, social security grants, and tax applications. Competences are both constituted and limited by law (whether written or unwritten).

Individual rights thus depend on the *institution of the rule of law*, that is:

- a distribution of public competences by way of a constitutional system of *checks and balances*; and
- *practical and effective fundamental rights* whose enforcement is disentangled from arbitrary decision-making by the government.

This will be further discussed in section 2.2 and throughout this book, notably when analysing the case law of the highest European courts.

2.1.3 Legal reasoning

If we understand law in terms of *legal conditions and legal effect*, the prominence of interpretation and argumentation becomes clear. This is connected with the possibility of contestation and the need for justification.

Legal reasoning is not just a matter of method, but first of all one of justification. It is not merely about heuristics but about legitimization. One could say that ‘solving’ a legal problem commences with heuristics, figuring out potential solutions. This will entail establishing the relevant facts (Peter hit Paula, who died), identifying potentially applicable legal norms (e.g. the criminal offence of negligent death, manslaughter, or murder), interpreting the facts in light of the norms (what if Paula is a cow, are the facts still relevant?) and interpreting the norms in light of the facts (what if Paula is a dangerous criminal who was on the verge of killing Peter?). After establishing and interpreting the facts in light of the norms and vice versa, a conclusion will present itself—based on the fact that if specific legal conditions apply, a specific legal effect will be attributed. Alternative solutions will also present themselves, as both the facts and the norms may be interpreted differently and the relevance or completeness of the facts as well as the identification of the applicable norm may be debatable. Sometimes, different norms with contradictory consequences

are applicable, requiring a higher-level decision on the priority of one over the other.

A crucial point, however, is whether a solution can be justified based on law. This is one of the pivotal functions of the law: to rein in arbitrary decisions based on prejudice or on the whimsical preferences of individual judges. The need to justify a decision constrains the 'solution space'. Justification thus affects the heuristics; it will generate self-censure as the judge knows they will have to justify their decision in legal terms. This *justification* can be portrayed as a syllogism:

Major: If a then b (legal norm)

Minor: a is the case (facts)

Conclusion: b (legal effect)

This scheme raises a number of questions that are best framed in terms of legal conditions and legal effect. As to the applicable legal conditions, the first question is which legal norm is relevant and how it relates to other relevant legal norms. Should the public prosecutor stick to murder or bank on manslaughter? To answer that question, the norm must be analysed in terms of the conditions it contains, for example, death of the victim, causation by an act or omission of the defendant, intent, and potential justification or excuse. The next question is whether these conditions are fulfilled, which requires an investigation of the facts, for instance asking which actions have been identified, which are missing, and which are relevant for the case at hand. These facts are historical events that must be *reconstructed* based on evidence such as witnesses, documents, forensic materials, and reporting, including *inferences* based on the available evidence, context, and common sense. The law of criminal procedure has strict requirements for what counts as lawful evidence (e.g. the police need a warrant for invasive investigation measures), and for the level of certainty that counts as proof that the offence has indeed taken place as charged.

This means that the second step (the minor) entails interpreting the facts in light of the relevant norm, while interpreting the relevant norm in terms of the facts.

This is a delicate operation that must be undertaken with great acuity, making sure that judgment is suspended until proof can be established

beyond reasonable doubt. Note that we are dealing with an example of criminal law, whereas the law of evidence and the burden of proof may differ in private law and administrative law. After deciding that the legal conditions apply, their legal effect must be established, which will be the final step in the 'solution' of a case. This will again demand interpretation. Criminal offences are usually threatened with a maximum punishment. This means that a court must weigh the seriousness of the offence and the culpability of the offender, taking into account numerous circumstances, before imposing a sanction. The fundamental legal principles of equality, fairness, and proportionality will require that similar circumstances will result in similar punishment, so the court will have to develop and sustain a policy to avoid arbitrary sentencing. This entails that the choice of punishment must be motivated.

A legal decision by a court and its anticipation by lawyers and citizens thus require a form of legal reasoning that explains and justifies the decision as lawful. This involves both more and less than logic, as the *ambiguity of human language* is part of the *protection* that law offers.

Application of legal conditions and legal effect is not a mechanistic affair. That is why legal reasoning is a matter of *argumentation* rather than logic, built on experience, expertise, and a salient acuity as to the many layers of interpretation that constitute legal judgment. Once judgment is given, legal effect is operational, based on the performative nature of legal decisions: if the accused is acquitted, she can legally ward off any punitive measures; if she is convicted, she can be imprisoned or fined accordingly.

The study of law is the study of legal conditions and legal effect. This entails an in-depth understanding of the sources of law and the arguments and lines of argumentation available for the justification of legal decision-making. In a sense, the study of law is about anticipating what a court will decide if confronted with the case at hand. As Oliver Wendell Holmes wrote: 'The prophecies of what the courts will do in fact, and nothing more pretentious, are what I mean by the law.'

What matters, however, is not merely the decision itself, but the legal reasoning that justifies it, as in the end the justification (what lawyers call the *ratio decidendi*) determines how a particular judgment will shape future case law.

2.2 What is Law in a Constitutional Democracy?

Law is closely related to politics (who decides the law?) and to morality (what content should prevail?). In many ways, law, morality, and politics are mutually constitutive. However, ‘in many ways’ does not mean ‘in any way’. In a viable constitutional democracy, law, morality, and politics cannot be related in an arbitrary fashion.

First, to some extent, *law shapes the playing field* for politics by the institution of legislative, administrative, and adjudicative powers that are both *enabled and constrained* by such institution. This refers to one of the core functions of the law: *the simultaneity of its instrumental and protective nature*. Law allows legal subjects, including the state, to act in law and to generate legal effect, but always conditioned by limitations that ensure, for example, legal certainty, proportionality, and transparency. Legal norms provide competences in a way that also protects interests, rights, and freedoms considered worthy of protection. Note that these interests, rights, and freedoms may be private interests, but *their protection is often deemed a public good*. Privacy, for instance, may be a *private interest* of an individual person, but its protection is also an important *public good* as it aims to sustain and protect the individual autonomy on which a vigilant democracy depends.

Second, to some extent, *law creates the level playing field that enables individuals, companies, and government agencies to act ethically*. The point of law is not to impose a specific morality on its constituency, but to provide the pre-conditions for developing an ethical stance and acting upon it. If companies are aware that data protection law prohibits cookie walls that force users to consent to privacy policies they would otherwise not consent to, they can develop other types of business models—knowing their competitors are forced to do the same.

Third, *law in a constitutional democracy constrains and enables both politics and morality in very specific ways*. Democracy is not the dictatorship of the majority but a system of checks and balances that requires a ruling majority to take into account that *democracy implies that minorities can become majorities*. This means that a ruling majority should not act in ways that pre-empt minorities from becoming a majority.

This goes back to what legal philosopher Ronald Dworkin considers the core of both democracy and the ‘rule of law’: governments should treat their

citizens as worthy of equal respect and concern. This grounds both the idea of one person one vote (representational democracy), and the imperative for majorities to respect individuals that are part of a minority (individual human rights).

2.2.1 Law, morality, and politics, and the nature of legal rules

One of the most famous legal philosophers of the twentieth century was Herbert Hart. In his seminal *The Concept of Law*, he explained the meaning of law in terms of three questions, aiming to set law apart from morality and politics.

The first question asks how law relates to and differs from orders backed by threats (commands).

Hart's answer regarding the *relationship* between law and commands is that modern positive law:

1. has teeth;
2. assumes state authority; and
3. depends on sovereignty but also constitutes it.

Hart's answer regarding the *difference* is that under the rule of law:

- legal norms apply to those who enact them (this distinguishes law from discipline or administration and 'rule of law' from a dictatorship);
- legal norms that confer legal powers to adjudicate or to legislate or to contract are not orders backed by threats (this relates to the difference between primary and secondary rules, see below under the third question);
- not all legal norms come into existence as explicit prescription (unwritten law, such as legal principles and customary law is not imposed by a legislature but confirmed by either the legislature or the courts); and
- sovereignty is not an apt description of law, even though law constitutes and limits it.

The second question is how legal obligation differs from and relates to moral obligation.

His answer regarding the *differences* is that modern law:

1. has teeth, whereas moral obligation is a matter of individual commitment; and
2. integrates primary rules with secondary rules that determine the validity of primary rules.

His answer regarding the *relationship* between legal and moral obligation is that law is not merely a matter of being forced to comply (e.g. the gunman situation). Having an obligation (in law as in morality) implies:

1. the existence of a standard;
2. its application to a particular person; which
3. may be against the interest of the person having the obligation.

The third question that Hart asks to clarify the nature of law inquires into the nature of legal rules. What are rules and to what extent is law an affair of rules?

Hart explains:

1. Legal rules are rules in the sense of *obligations*, not rules in the sense of regularities. The mere fact that most people violate a traffic rule does not stop it from being a legal rule.
2. He notes that rules are observed from an *internal point of view*, they assume a sense of obligation. Even when one violates a legal rule, one supposedly remains committed to the obligation to comply. In a sense this is the core difference between law and force: the possibility to disobey the law is constitutive of the law; validity does not depend on brute force in itself.

This raises the question of what determines the validity of legal rules. Hart's brilliant answer was that this is decided by law itself, in a highly distinct way.

Legal rules, he proposes, come in two types: primary and secondary rules.

- Primary rules are *regulative rules*, they 'regulate' our interactions by imposing a prescription or a prohibition (e.g. 'you shall not kill').
- Secondary rules are *constitutive rules* that determine the validity of primary rules and the legal effect of violation.

Secondary rules *confer powers*, for example, ‘Parliament decides on criminalization,’ ‘if you kill you will be punished with . . .,’ ‘to be legally married the marriage must be inscribed in the civil registry.’ Hart argues that the difference between primary and secondary legal rules is typical for modern positive law, as it allows a court to authoritatively determine the *validity* and thus the *applicability* of legal norms without resorting to either regularity or brute force.

This highlights the systemic and architectural nature of positive law, which consists of a complex, coherent system of primary rules that clarify what is expected, supported by secondary rules that allow one to test whether a primary rule is indeed valid.

2.2.2 Legal certainty, justice, instrumentality

We end with the concept of law that was introduced in the beginning of this chapter, based on the work of Radbruch. The reason for selecting Radbruch is that he pins down three goals that law must serve, without ignoring the fact that in concrete cases these goals are often incompatible. He withstands the temptation to reduce two of these goals to sub-goals of one and thus to resolve the tension between them. Instead, he highlights the importance of nurturing this tension, sustaining it, and thus challenging lawyers to continuously reinvent the right balance or trade-off, without thereby discarding any one of the three goals as being overruled. This accords with a difference of opinion between Dworkin and Hart about the nature of law. Whereas Hart initially claimed that modern law can be characterized as a system of legal rules, which are either applicable or not, Dworkin argued that the decision as to which legal rule applies and how it must be interpreted in concrete cases involves an important role for *legal principles*. Other than rules, Dworkin said, principles do not follow the binary applicability of rules. Principles have a certain *weight*, depending on what is at stake in the case at hand. In the case of competing rules, either one will ‘win’. In the case of competing principles, both can be relevant and both can inform the decision (notably the decision as to which rule is valid), though their impact on the decision may vary. For Radbruch, who served as Minister of Justice in the Weimar Republic in the 1930s, the tension between justice, legal certainty, and instrumentality in law was not mere intellectual nit-picking. The rise of Nazism and the role of law as an instrument of genocide challenged the balance between law’s fundamental goals.

Before elaborating on that, I will first clarify how *the goals of the law* can be understood.

- *Legal certainty* refers to the need to provide a *foreseeable response to one's actions, in order to create societal trust*. For Radbruch, this refers to law's 'positivity' or 'positiveness', that is to the fact that law is 'posited' by a legislature (and by the courts that decide its correct interpretation). The legal power to 'posit' the law is based on what Hart termed a set of secondary rules that determine the validity of legal norms. Though Radbruch was not a positivist in the sense that he only cared about the formal validity of legal norms, he attached particular importance to the 'positivity' of law and the legal certainty it provides. He explains that precisely because we may not agree about what moral duties we have, law provides a measure of certainty about the legal rights we have and the legal obligations we should comply with. Legal certainty is also connected with the notion of equality before the law; it is the opposite of arbitrary, discriminatory, or explicitly unjust exercise of state authority. The goal of legal certainty does not necessarily overrule the other two goals; if so, their ranking would collapse into a positivism that separates law entirely from both morality and politics—thus turning it into an unresponsive and mechanical form of administration.
- *Justice* refers to *treating equal cases equally, and unequal cases unequally to the extent of their inequality*. This is directly connected with legal certainty as this should enable people to plan ahead, being capable of anticipating how their actions will be 'read' by the law and responded to by the government. That also goes for how the government responds to actions by others that concern us (criminal offences, breach of contract, invasion of privacy by a private company). But, as Dworkin argued, justice is *more than mere consistency*; it is rather about the *integrity* of the totality of legal rules, principles, and policies, ensuring that each decision is taken in accordance with the *implied philosophy that grounds the law*. Justice as fairness concerns two types of equality (as described by Aristotle): distributive and corrective. Distributive justice means that everyone should be treated in the same way, to the extent that similar conditions apply. Corrective justice means that punishment should be proportional to the seriousness of the crime and compensation proportional to the damage suffered. Clearly, both types of justice are related, as the determination of punishment or compensation must be aligned with the relevant crime

or tort as well as with the legal response with regard to similar crimes or torts. As indicated, Radbruch highlights that the goal of justice does not necessarily overrule the goals of legal certainty and instrumentality. If so, law would collapse into a moralistic enterprise.

- *Instrumentality* refers to the fact that law is *an instrument to achieve a variety of goals* that are in part external to its own operations. These goals play out at the level of politics (legislation), where law is a policy instrument and at the level of individual legal subjects (including companies) who will use the law to strategically further and protect their own interests (private law) and their rights and freedoms (private law, criminal law, and legal remedies afforded in administrative law). Since the rise of independent courts in sixteenth to eighteenth-century Europe, law and politics have struck a historic bargain: *law does not interfere in politics (where goals are to be determined by democratic legislation), while politics remains under the Rule of Law*. This means that though law is instrumental in achieving the goals that a democratic legislature determines, it has its own values and goals that will constrain the ‘solution space’ of political goal-setting and its execution. Once again, Radbruch reminds us that the goal of instrumentality does not necessarily mean that law’s expediency will overrule justice and legal certainty. If so, ‘rule of law’ would collapse into arbitrary rule *by law*, and law’s instrumentality would reduce to instrumentalism.

The three goals of the law are constitutive of law and the ‘rule of law’ (they determine what *counts as law*), but they are also antinomian (they may be incompatible in concrete cases). When the Second World War ended, Radbruch wrote a brief text to explain how his antinomian goals relate to Nazi rule. The title of his text was: *5 Minutes of Legal Philosophy*. He targets some of the maxims that were typical for the way that law was instrumentalized by Nazi Germany. First, the maxim of ‘an order is an order’ and ‘a law is a law’. He frames this as the equation of law with power. Second, the maxim of ‘law is what benefits the people’ and ‘whatever state authorities deem to be of benefit to the people is law’. This results in framing the private benefit of those in power as equivalent with public benefit. Instead of these populist maxims, he presents ‘law as the will to justice’ and ‘equality before the law’. He reiterates that law is determined by the antinomian goals of instrumentality, justice, and legal certainty, and adds that laws that do not even aim for justice do not merely forsake their validity within the system, but *must be denied their legal character*. This demonstrates the priority of fundamental principles of law that preclude mistaking brute force or arbitrary rule *by law* for the rule of law.

We can sum up this chapter by stating that in a constitutional democracy, legal rules that confer powers simultaneously restrict them; they provide functionality *in a way that* provides protection, thus serving the double instrumentality of the law as a tool of both government and protection.

References

On the concept of law

- Hart, H.L.A. 1994. *The Concept of Law*. Oxford: Clarendon Press.
- Hildebrandt, Mireille. 2015b. 'Radbruch's Rechtsstaat and Schmitt's Legal Order: Legalism, Legality, and the Institution of Law'. *Critical Analysis of Law* 2 (1): [1].
- Holmes, Oliver Wendell. 1997. 'The Path of the Law'. *Harvard Law Review* 110: 991–1009.
- Radbruch, Gustav. 2014. 'Legal Philosophy'. In *The Legal Philosophies of Lask, Radbruch and Dabin*, translated by Kurt Wilk. Introduction by Edwin W. Patterson, 44–224. Twentieth Century Legal Philosophy Series 4. Boston and London: Harvard University Press.
- Wesel, Uwe. 1985. *Frühformen Des Rechts in Vorstaatlichen Gesellschaften. Umriss einer Frühgeschichte Des Rechts Bei Sammlern Und Jägern Und Akephalen Ackerbauern Und Hirten*. Frankfurt am Main: Suhrkamp.

On the sources of law

- Hage, Jaap, Antonia Waltermann, and Bram Akkermans. 2017. *Introduction to Law*. 2nd ed. New York: Springer, chapter 1.

On speech act theory

- Austin, J.L. 1975. *How to Do Things with Words*. 2nd ed. Boston: Harvard University Press.
- Searle, John. 1995. *The Construction of Social Reality*. New York: The Free Press.

On speech act theory in law

- Hildebrandt, Mireille. 2015a. *Smart Technologies and the End(s) of Law. Novel Entanglements of Law and Technology*. Cheltenham: Edward Elgar.
- MacCormick, Neil, and Ota Weinberger. 1986. *An Institutional Theory of Law: New Approaches to Legal Positivism*. Dordrecht, Boston and Hingham: D. Reidel Pub.

Co.; sold and distributed in the United States and Canada by Kluwer Academic Publishers.

On legal reasoning

Dickson, Julie. 2010. 'Interpretation and Coherence in Legal Reasoning.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward Zalta. <http://plato.stanford.edu/archives/spr2010/entries/legal-reas-interpret/>.

Dworkin, Ronald. 1982. 'Law as Interpretation.' *Texas Law Review* 60 (2): 527–50.

Gigerenzer, Gerd, and Christoph Engel, eds. 2006. *Heuristics and the Law*. 1st ed. Cambridge: The MIT Press.

Hage, Jaap, Antonia Waltermann, and Bram Akkermans, eds. 2017. *Introduction to Law*. 2nd ed. New York: Springer, chapter 2.

On legal rights

Edmundson, William A. 2012. *An Introduction to Rights*. 2nd ed. Cambridge: Cambridge University Press.

Hohfeld, Wesley Newcomb. 1964. *Fundamental Legal Conceptions, as Applied in Judicial Reasoning*, edited by Walter Wheeler Cook with a new foreword by Arthur L. Corbin. New Haven: Yale University Press.

On the influence of the Snowden revelation on EU data protection law

Granger, M.-P., and K. Irion. 2014. 'The Court of Justice and the Data Retention Directive in Digital Rights Ireland: Telling Off the EU Legislator and Teaching a Lesson in Privacy and Data Protection.' *European Law Review* 39 (6): 834–850.

On law, democracy, and the rule of law

Dworkin, Ronald. 1991. *Law's Empire*. Glasgow: Fontana.

Radbruch, Gustav. 2006. 'Five Minutes of Legal Philosophy (1945)'. *Oxford Journal of Legal Studies* 26 (1): 13–15. <https://doi.org/10.1093/ojls/gqi042>.

Taekema, Sanne. 2013. 'The Procedural Rule of Law: Examining Waldron's Argument on Dignity and Agency.' In *Annual Review of Law and Ethics. The Rule of Law-Principle* 21:133–46, edited by B.S. Byrd, J. Hruschka, and J.C. Joerden. Berlin: Dunckler and Humblot. <http://papers.ssrn.com/abstract=2391228>.

Waldron, Jeremy. 2008. 'The Concept and the Rule of Law.' *Georgia Law Review* 43 (1): 1.

3

Domains of Law: Private, Public, and Criminal Law

Computer science can be divided into a plethora of different subdisciplines, while division may depend on whether one comes from, for example, electrical or electronic engineering, from mathematics, software engineering, statistics, cognitive science, or machine learning. Law and the study of law is most often divided in three major domains: private, public, and criminal law. These domains have their own principles, own vocabularies, and structures, geared to the type of relationships they concern. For instance, when relationships are vertical, as in public and criminal law, different principles apply than when relationships are considered horizontal, as in private law.

This chapter will first explain how these domains differ, based on a set of conceptual distinctions. This provides the foundations for the subsequent introduction of the core structure, vocabulary, and underlying principles of each domain. This is pivotal for a proper understanding of more specified domains such as data protection law, cybercrime, and copyright that comprise the second part of the book.

3.1 Private, Public, and Criminal Law: Conceptual Distinctions

If we ask the question of ‘what law does’, the answer is as simple as it is complex: law creates legal effect. The complexity resides in how this is done, even though here again the answer seems simple: this depends on the applicable legal conditions. To identify the relevant legal conditions, we must search the sources of law (a concept with a very specific meaning, as explained in section 2.1.1).

In section 2.2.1, we introduced law as a system of legal norms, notably as a combination of primary and secondary legal rules. These rules form a complex architecture with multiple dimensions (e.g. local, national, international, and supranational rules; general and more specific rules; prior and posterior

rules; legislation and case law that enact and interpret rules; and principles that are derived from the implied philosophy of positive law).

To clarify the difference between private, public, and criminal law, we will add a complementary perspective to frame the law.

Next to describing law as a unity of primary and secondary rules and underlying principles, we will picture law as a system of *legal relations* between *legal subjects*, with regard to *legal objects*.

3.1.1 Absolute rights and relative rights

Property rights, such as ownership, are often described in terms of the relationship between a legal subject (e.g. a natural person or a corporation) and a legal object (e.g. a house or a receivable), stating that the subject has a right in the object. To better understand what this means we will describe property rights in terms of the relationships between legal subjects, *with regard to* a legal object.




Figure 3.1 Absolute and relative rights

In Figure 3.1 we can see that a legal subject with an absolute right in a legal object imposes a duty of non-interference for ALL other legal subjects, with respect to her right in this legal object. Property rights are thus *absolute*

rights in this particular sense: all others must refrain from interfering with these rights. Absolute, here, does not refer to unlimited. Property rights may be limited by, for example, a prohibition to abuse the right, or by human rights. For instance, if I own a house and rent it out to someone, my property right in the house is not unlimited; I cannot enter the house at will, because this may violate the right to privacy of the person who rents the house.

Though not unlimited,¹ absolute rights can in principle be enforced against all legal subjects.

In Figure 3.1 we can also see that a relative right only plays out between a restricted set of legal subjects. For instance, a contract usually generates two legal effects, that is, the legal obligations to which the contract commits the parties. In the case of a contract of sale, one party will have to pay the agreed price, the other party will have to deliver the agreed good or service. Both parties have the legal right that the other party complies with her legal obligation. But, other than in the case of property rights, these rights only apply to the relevant party to the contract. There is no duty for other legal subjects to comply with the agreed legal obligations.

Relative rights can only be enforced against specific legal subjects.

3.1.2 Private law and public law

Many attempts have been made to find conclusive criteria to distinguish private and public law. For instance:

Whenever the government is involved, we are in the domain of public law.

This would mean that if a government agency buys pencils, the contract would be ruled by public law. The seller of the pencils may object that this exempts the government from the ‘rule of law’, as this would exempt it from, for example,

¹ Note that in human rights law the concept of ‘absolute rights’ has a different meaning, referring to rights that cannot be limited in any way (e.g. the right against inhuman and degrading treatment), as opposed to rights that can be limited when confronted with competing human rights or a public interest (e.g. the right to privacy, which is in this sense *not* an absolute right as it may be overruled by e.g. freedom of information or public health).

the duty to pay compensation in the case of breach of contract. Therefore, in constitutional democracies, this criterion is not conclusive. When the government buys pencils, private law applies.

Another criterion suggests that:

Whenever the public interest is involved, we are in the domain of public law.

Though this sounds plausible, it seems to be in the public interest that parties to a contract are bound to comply with the obligations to which they have committed. If the criterion of the public interest is applied, this would mean that such compliance is part of public law. This is clearly not the case. In constitutional democracies, the inverse does pertain; the government is bound to always act in the public interest.

Finally, yet another criterion proposes that:

Public law entails that the enforcement initiative is with the government.

In private law, enforcement is left to private parties. They can go to court or, for example, involve a bailiff, but the government will not take the initiative to enforce compliance with a contract. This is connected with the idea that, in private law, parties are autonomous as to the content of a contract but also with regard to how they respond to defaulting by the opposite party. This might lead one to conclude that when public law is at stake, the enforcement initiative is with the government. However, in administrative law, citizens can take the initiative to object or appeal against a decision made by the administration. This does not turn the *legal remedies* of citizens into private law. A legal remedy is the legal power to contest a decision or action in a court of law, thus e.g. achieving annulment or avoidance of the decision, compensation, or injunctive relief (a court order that unlawful conduct is terminated).

Instead of trying to develop criteria, we can resort to a simple inventory.

Public law consists of:

1. constitutional law;
2. administrative law; and
3. international public law.

We could say that in these subdomains the government acts ‘as such’, that is, in its capacity as a public authority, under the rule of the legality principle. As soon as the government acts in its capacity as a private party, private law will apply.

It follows that private law is that part of law where the government ‘as such’ does not play a role. We must remind ourselves, however, that private law is an artificial construction, just like public law and criminal law. It has been instituted by the legislature and will be enforced and fine-tuned by the courts. In that sense, its construction is based on the attribution of legal powers to legislate and adjudicate and thus depends on public law (the constitution). It is tempting to frame private law as ‘given’ or ‘natural’ law, as if it is merely the written articulation of an existing unwritten private law. This temptation must be resisted as it hides the fact that legislators and courts make many choices when deciding on the content of private law, while these choices basically constitute and regulate economic markets. It is not the market that dictates the power of law, but the law that ‘affords’ a specific type of economic market (that may in turn enhance or diminish legal protection).

Above, we observed that in a constitutional democracy, the government must always act in the public interest. This raises the question of the *purpose of private law*.

- Individual citizens are not by default required to act in the public interest, instead private law gives them the legal tools to act strategically in their own interest. This is related to the idea of *individual autonomy* that seems to be the hallmark of private law.
- Private law provides legal norms meant to create a private sphere where companies, consumers, employees and employers, service providers, and users are in principle free to conduct a business, to conclude contracts, and to navigate their personal, social, and institutional environments as they wish. Private law thus aims to create and sustain *societal trust*, based on legal certainty.

For instance, when one buys a house, one can in principle be sure that the owner can be forced to deliver the house after the price has been paid. And, when a person wrongfully causes damage to another, the other must in principle be sure that the tortfeasor can be forced to compensate for the damage. I qualify these general rules by inserting ‘in principle’, because exceptions apply.

Next to individual autonomy and societal trust, private law is also about *fairness*. For instance, in shaping economic markets, private law not only ensures that agreed prices must be paid and goods delivered, but subdomains such as, for example, consumer law, non-discrimination law, competition law, product liability, and unfair contract terms law have dimensions of fairness.

These subdomains aim to compensate the lack of bargaining power of weaker parties (e.g. consumers) or to protect specified groups against unjustified discrimination. This demonstrates that private law can be restricted, for instance by constitutional limitations and international human rights law, but also by administrative law. House owners have full disposition of their property, but administrative law may restrict their competence to renovate the house, for instance based on safety requirements that are part of administrative law.

The *purpose of public law* is the *public interest* or benefit, in Latin the *res publica*, which resulted in the idea of a Republic. The public interest has a very broad meaning and basically requires an assessment of a diversity of public interests, which may be incompatible in concrete situations. Public law concerns, for instance, safety and security, welfare, public health, care for the elderly, public education, public traffic management, full employment, public housing, etc. Public law, notably administrative law, is restricted by the *legality principle* that requires a legal basis for all acts and decisions of the government. This legal basis can be very general if the actions or decisions do not entail negative consequences, but if negative implications can be expected, the legal basis must be specific in what it allows, for what purpose, and under which precise conditions. If the government, for instance, wants to disown a person to enable the construction of a new public road, highly specific conditions apply, and the person must be compensated.


The *legality principle* demands that the government always acts within the limits set by the written or unwritten Constitution (that attributes powers to the government). This is directly related to the fact that under the 'rule of law', legal norms are both constitutive and limitative of the legal powers they attribute (section 2.1.2.1).

On top of that, international human rights law and other treaties to which a state is bound, will further restrict the powers of the government.

Some legal norms are *mandatory*, which means that they cannot be overruled by contractual or other norms. In public law, most of the norms are mandatory, both when addressing citizens (e.g. prescribing with what conditions

they must comply to obtain a building permit) and when addressing the government (e.g. prescribing under what conditions a municipality must grant the building permit). In private law, many legal norms are *default*, especially in the domain of contract law, meaning that such norms only apply if parties have not agreed otherwise.

Returning to the perspective of law as a complex architecture of legal relationships between legal subjects, we can depict public law as follows:



Public law creates a legal relationship of public authority with all citizens, who are thus all placed in legal relationship with each other.

[Please note that all legal subjects should be connected to all legal subjects with grey lines]

Figure 3.2 Public law as an architecture of legal relationships

Figure 3.2 shows that public law is based on a relationship between each individual legal subject and the state. This is an example of distributive equality, meaning that all citizens are at equal distance from the state, being entitled to equal respect and concern (section 2.2.2). This creates a specific type of equality amongst citizens, who—even if they remain strangers to each other—share an equivalent relationship to the same state. Based on this relationship to the state, citizens can develop legitimate mutual expectations, knowing that the state can enforce such expectations if they are ‘covered’ by legal norms.

3.1.3 Private law and criminal law

Let’s look at a typical exam multiple choice question on the topic of the difference between private and criminal law:

- I. If downloading of illegally provided content is unlawful, it is necessarily punishable.

- II. If someone commits a crime the victim can initiate proceedings.
- a. I and II are both correct.
 - b. I is correct, II is not correct.
 - c. I and II are both incorrect.
 - d. I is not correct, II is correct.

This question aims to test a proper understanding of the difference between an action being unlawful and an action being punishable. Clearly, for an action to be punishable, it must be unlawful. It would be against the criminal law legality principle to punish a person if such action was not clearly defined as being unlawful. But the criminal law legality principle demands more than that. It requires that an action can only be punished if it was *clearly defined* as a criminal offence at the time of committing the offence. If not, the action does not qualify as an offence. This means that most unlawful actions are not criminal offences.




Figure 3.3 Unlawful and criminal conduct

Criminal conduct is a subset of unlawful conduct. Not everything that is unlawful or illegal is punishable. Whether something is unlawful and/or criminalized conduct depends on whether it has been defined as such in the *objective law*, which also defines any *subjective rights* a person may have.

- *Objective law* is the unity of primary and secondary rules and the implied principles of law that is valid within a specific jurisdiction.
- A *subjective right* is a right attributed to a legal subject by the objective law; a legal right therefore depends on the objective law that grants it.

In Chapters 1 and 2, notably in sections 1.4 and 2.2.1, we discussed the concept of positive law, which is close to that of objective law. Positive law is the law that is in force in a specific jurisdiction. The emphasis here is on the ‘posited’ or artificial nature of modern law. Objective law refers to the same, but juxtaposes the assemblage of rules and principles to the set of individual rights, highlighting the fact that subjective rights depend on objective law.

Who or what is a *legal subject* is not given, but depends on objective or positive law. A legal subject (a natural person or a legal person) is an entity capable of acting in law, bearing legal rights and legal obligations in relation to other legal subjects.

In contemporary positive law we distinguish between two types of legal subjects: a *natural person* or a *legal person*. Positive law decides what entities have legal personhood. Think of corporations, municipalities, or the state itself. These all have standing in law, they can conclude contracts and be held liable. In principle, positive law could attribute legal personhood to animals or robots. We will return to this point in Chapter 9.

A *legal object* is an entity that is the object of a specific legal relationship between legal subjects.

Think of a legal good such as an intellectual property right, a tangible, or a specific obligation. If I conclude a contract with one of my students, to sell her a book, the book is the legal object. More precisely, the property right of ownership is the legal object that will be transferred.

3.2 Private Law

Private law can be subdivided into, for example: family law (marriage, inheritance); contract law (general, specific); property law (transfer of ownership); and tort law (general, specific). In this monograph we will not discuss family law, but focus on contract, property, and tort law.

Private law contains the legal norms that regulate relationships between legal subjects at a *horizontal* level, thus excluding the government acting ‘as such’. Horizontal does not mean that legal subjects are equal in the sense of having the same economic or other power. It means that they are formally considered as equal, capable of determining their own position in law; they

can, for example, conclude contracts, sell their property, and be held liable for unlawful wrongdoing. As indicated above, private law respects the autonomy of individual persons but also contributes to such autonomy by enabling them to act strategically within the private sphere, as long as they act within the bounds of the law. Other ways of contributing to individual autonomy can be detected in, for example, consumer law and competition law, which aim to compensate weaker parties with less bargaining power. We will now discuss property law, contract law, and tort law.

3.2.1 Property law: transfer of movables

Alice has a book, Bob has a house, Eaves has a wonderful surname (Dropping). Legally speaking, the question is what ‘has’ means. Does Alice own the book, does Bob rent the house, can Eaves sell her name? These questions bring us into the heart of property law. A surname cannot be sold, it is not a property—even though it clearly belongs to Eaves and not to another. But does ‘having’ a house in the sense of ‘renting’ turn the house into the property of Bob? Renting a house means that one has the right to live in the house of the owner, based on the freedom of the owner to rent out the house. In most jurisdictions, there is a subdomain of private law dedicated to the renting of real estate (including protection of those who rent a home against arbitrary decisions of the owner).

What if Alice has borrowed the book and sells it to Bob? In this case Alice first *held* the book for another (the owner from whom she borrowed). When she sold it to Bob, she *possessed* the book (this means that from that moment onwards she is holding the book for herself). *Though she possessed the book, she did not own it*, because ownership would imply possession with right. So, Alice first held the book, then possessed it, but never owned it. The question is whether Bob owns the book, after buying it from Alice.

This is a question concerning transfer, possession, and ownership of movables or tangibles. This question must be answered with regard to a specific jurisdiction, because private law is not the same in each country. Because *positive law is posited*, different legislatures and courts can ‘posit’ different rules about transfer, possession, and ownership of movables. Let’s take an example from the Netherlands Civil Code (NCC), Article 3:84 NCC (Requirements for a transfer):

1. The transfer of property requires a delivery pursuant to a valid legal basis by the person with power of disposition over that property.

This means that the law requires that three legal conditions are fulfilled to achieve the *legal effect of a transfer of movable property*:

1. delivery;
2. valid legal basis or title; and
3. power of disposition.

Note that these conditions are cumulative; each condition must be fulfilled. We will now check whether the property of the book is transferred. For a *delivery* it would be enough that Alice hands over the book to Bob. If Alice has concluded a valid contract of sale with Bob, there is a *valid legal basis* or *title*. However, since Alice is not the owner, she lacks *power of disposition*. One of the crucial legal powers that the law attributes to ownership, is the power of disposition, or the freedom to share, sell, give, or even destroy the object of ownership. It seems to be that Alice cannot transfer property, because she has no power of disposition.

Now, check Article 3:86 NCC (Lack of power of disposition)

1. A transfer of a movable thing (...) by an alienator without power of disposition is nevertheless valid if the transfer was not performed gratuitously and the acquiring party acted in good faith.

Here we see that the legal effect that Bob wants to achieve, transfer of the book, can be reached despite the fact that Alice (the alienator) has no power of disposition.

Without power of disposition, a valid transfer of a good can nevertheless be achieved, if the following legal conditions are fulfilled:

1. movable;
2. transfer not for free; and
3. good faith of the acquiring party.

Again, note that the conditions are cumulative. The book is a movable, and the sale presumes that a price has to be paid so the transfer is not for free. The final condition concerns the *good faith* of Bob. If he *knew (or should have known)* that Alice borrowed the book, the ownership will not be transferred. If he was not (and should not have been) aware of that, he will become the owner. Note that this implies that Alice has transferred something

that she did not have herself: the *ownership* of the book. This is an exception to an important written or unwritten rule that is part of private law in most jurisdictions: *nemo dat quod non habet* (no one gives what they do not have). The exception is based on the need to ensure trust in economic relationships. In this case, the buyer is protected because they should be able to assume that a person who possesses a movable is the owner. The idea is that this smooths day-to-day economic transactions, which would become cumbersome if one first has to figure out whether the person who is selling is actually the owner.

Now, what if Alice stole the book? The Netherlands legislature wants to contribute to a transparent marketplace, where people can trust that items clearly possessed by the seller will become their property (if they have no reason to believe the seller is not the owner). However, the legislature does not want to reward the theft of goods. Therefore, we have Article 3:86 NCC:

3. The owner of a movable thing who has lost possession of it because it was stolen from him may, in spite of the previous paragraphs, always claim his property back from every possessor within three years after the theft, unless:
 - a. the stolen object has been acquired by a natural person who, when he acquired it, did not act in the pursuance of his practice or business, and who had received it from an alienator who sells these or similar objects regularly to the public making use of a business premises destined for that purpose and who acted, when he passed the stolen object to the acquiring party, in the conduct of his practice or business, yet not as an auctioneer;
 - b. or the stolen object concerns money or negotiable documents for a claim to order or to bearer.

Again, we first identify the relevant legal effect. In this case, the legal effect concerns the owner of the stolen book. They can claim back (revindicate) their property, if the following legal conditions apply:

1. less than three years have passed since the book was stolen; and
- 2.a Bob is not a natural person; or
- 2.b Bob is acting in the context of his practice or business; or
- 2.c Alice does not regularly sell second-hand books to the public in business premises destined for that purpose; or
- 2.d Alice did not pass the book to Bob in the conduct of her business; or

- 2.e Alice passed the book to Bob as an auctioneer; or
3. the book does not count as money or negotiable documents for a claim to order or to bearer.

Note that paragraph 3 of Article 3:86 consists of one positive condition, followed by a series of negative conditions under Article 3.a and 3.b. These conditions are not cumulative, if any of them applies, the legal effect (of revindication) cannot be attributed.

If we assume that Alice does not have a business of selling things like books, the original owner will be able to revindicate her book from Bob within three years of the theft.

This piece of legislation nicely demonstrates how the law protects both the interests of different private parties and the general interest of legal certainty and trust in business transactions.

3.2.2 Contract law and property law: sale and transfer of real estate

Can Bob sell the house he rents to Eaves? The first question we need to confront is whether selling is a matter of property law or contract law. When we ask if Bob can sell a house, we are inquiring whether he can conclude a contract with Eaves about the sale of the house. This is a question of contract law, which is a subdomain of the law of obligations. Let us check the legal definition of a contract in Article 6:213 NCC on the 'Definition of an obligatory agreement'

1. An agreement (contract) in the sense of this Section is a more-sided (multilateral) juridical act under which one or more parties have subjected themselves to an obligation towards one or more other parties.

The legal effect consists of a valid obligatory agreement, which comes into existence if the following legal conditions are fulfilled:

1. a more-sided (multilateral) juridical act,
2. by which one or more parties subject themselves,
3. to an obligation towards one or more other parties.

These conditions are *cumulative*, the legal effect only occurs if all three conditions apply. This raises the question of what is meant by a ‘juridical act’, defined in Article 3:33 NCC on ‘Intention and declaration’:

1. A juridical act requires the will (intention) of the acting person to establish a specific legal effect, which will (intention) has to be expressed through a statement of the acting person.

The legal effect is the existence of a valid juridical act, and the legal conditions are:

1. the will or intention of the acting person to achieve a specific legal effect;
2. the will has been expressed through a statement of the acting person.

So, in principle, if Bob and Eaves expressed their intention to transfer property (Bob) and to pay a price, which means to transfer money (Eaves), they have concluded a valid agreement.

In this case the agreement is a contract of sale, which generates two legal obligations:

1. the buyer must pay the agreed price; and
2. the seller must transfer property.

The second question we confront, if Bob actually manages to sell the house to Eaves, is whether he can transfer the property of the house to Eaves. This is a matter of property law, just like the transfer of movable property. As we have seen above, Article 3:84 NCC requires the power of disposition to transfer property. Since Bob is not the owner, he lacks that power. In the case of real estate, the exception for the transfer of movables does not apply. In principle, Bob can, therefore, sell the house but he cannot transfer the property. This entails that, though the contract of sale has been concluded, he will not be able to fulfil his legal obligation to deliver the house, and thus Bob will be ‘in breach of contract’.

What if Eaves had counted on the transfer of property and suffers damage due to Bob’s incapacity to deliver the house? We now check Article 6:74 NCC on ‘Requirements for a compensation for damages’:

1. Every imperfection in the compliance with an obligation is a non-performance of the debtor and makes him liable for the damage which the creditor suffers as a result, unless the non-performance is not attributable to the debtor.

The legal effect here is the liability for damages, the legal conditions are:

1. there is non-performance of the debtor (the one who did not perform),
2. as a result of an imperfection in the compliance with their obligation,
3. the non-performance is attributable to the debtor.

So, in principle: Bob can sell the house, but he cannot deliver it. This is a salient example of the difference between absolute and relative rights, as discussed above: absolute rights are rights with regard to a good, that can be sustained against everybody; they create a duty of non-interference for all others. This is why their *publicity* is crucial; anybody must be able to check the relevant legal effect. For tangibles, *possession is the default form of publicity*: when someone holds a tangible for themselves, we can assume it is theirs. For real estate, we have a *public registry* where people can check who has a property right; since nobody carries their real estate with them, possession does not mean much. Relative or personal rights are rights that can only be sustained against a specific person or persons; since third parties cannot derive rights from relative rights that do not concern them, and have no obligations to respect relative rights that do not concern them, they have no need to know, so by default such rights are not registered.

In private law there is only a limited set of absolute rights, which is again related to the fact that everybody has an obligation not to interfere with these rights.

In most jurisdictions this set of property rights consists of: ownership, freehold, leasehold, servitude, right of superficies, apartment right, usufruct, pledge and mortgage, and intellectual property rights (e.g. copyright, patent). We speak of this as a closed system of property rights, as new rights cannot be created at will by individual legal subjects, even if they would agree. If the reader wants to know more about the content of these rights (their legal conditions and the ensuing legal effect), they are advised to check the relevant literature (under the references of this chapter).

Relative rights usually form an open system, where people can create new rights by way of contract (next to the contract of sale or rent or employment).

Non-contractual relative rights are: tort (e.g. violation of privacy), undue performance, and, unjustified enrichment. In this chapter we will deal with the most important type of non-contractual right, based on tort liability.

3.2.3 Tort liability

To understand tort law, we shall now distinguish a *juridical act* from a *juridical fact*. As discussed above, a juridical act is an action that aims for the legal effect the law attributes, for instance the validity of a contract, the validity of a will, or legislation that is in force. Often, however, the law attributes legal effect even if this was not intended. A juridical fact is an occurrence, status, or act that is legally relevant because the law attributes legal effect, irrespective of intent, for instance birth (attribution of legal subjectivity), death (inheritance), and tort (liable to an injunction and/or compensation of damage).

To understand the complexities of tort law I will discuss the famous Dutch ‘cellar hatch case’, which was decided in 1965 by the Netherlands Supreme Court.² The facts of the case are as follows. In 1961, in Cafe De Munt at Singel 522 in Amsterdam, Duchateau goes to the loo and falls into a cellar that was open, because ‘Sjouwerman’ (working for Coca Cola) was busy putting drinks in the cellar and left the hatch open. Duchateau suffered serious harm and sued Coca Cola for the damages. He did not sue Sjouwerman himself, because Coca Cola had deeper pockets and, by default, an employer is liable for damage caused by one of its employees, if it has been caused during normal working operations.

The legal question at stake was whether Sjouwerman should have taken into account the fact that people may not be as cautious as required to prevent the accident. This is a crucial question as the default of private law is that everyone carries their own damages. Some people may have bad luck due to disease, an accident or whatever, but unless the law makes an exception, such bad luck cannot be charged to another. One of these exceptions is a tort.

A tort is defined in Article 6:162 NCC:

1. A person who commits a tortious act (unlawful act) against another person that can be attributed to him, must repair the damage that this other person has suffered as a result thereof.

The legal effect is a legal obligation ‘to repair the damage’ (to pay compensation), and the cumulative legal conditions are:

1. a person has committed a tortious act (an unlawful act),
2. that can be attributed to him (attribution of act to tortfeasor),

² Netherlands Supreme Court, 5 November 1965, NJ 1966, 136.

3. against another person, who suffered damage (damage),
4. the damage is the result of the tortious act (causality between act and damage).

To decide whether an act (including an omission) *counts as* a tortious act, the second paragraph of Article 6:162 NCC stipulates:

2. As a tortious act is regarded a violation of someone else's right (entitlement) and an act or omission in violation of a duty imposed by law or of what according to unwritten law has to be regarded as proper social conduct, always as far as there was no justification for this behaviour.

To qualify as a tortious act, three alternative conditions and one cumulative condition apply:

1. the act was a violation of another's right; or
2. the act was a violation of a legal duty; or
3. the act violates an unwritten legal duty; and
4. there was no justification for the act.

To decide whether the tort can be attributed to the tortfeasor, the third paragraph of Article 6:162 NCC stipulates:

3. A tortious act can be attributed to the tortfeasor if it results from his fault or from a cause for which he is accountable by virtue of law or generally accepted principles (common opinion).

To qualify as an act 'of the tortfeasor', the following alternative conditions apply:

1. the act results from his fault (culpability); or
2. the act results from a cause for which he is accountably by virtue of law or generally accepted principles (risk liability).

Beyond these two types of attribution (of the act to the tortfeasor), most jurisdictions distinguish between: *fault liability* (culpability), *vicarious liability* (where another is liable for a tortious act, e.g. in the case of an employer being liable for the tortious acts of their employees), and *strict liability* (e.g. of an owner for the animal they keep or the building they own).

Some jurisdictions also discriminate between *risk liability* as an inversion of the burden of proof, meaning that exculpation is possible, and strict liability

where exculpation is not possible. This relates to the attribution of *causality*: if there is damage and the damage could reasonably likely have been caused by the tortious act, causality is assumed. In the case of risk liability, the tortfeasor can still prove they did not ‘cause’ the damage (often termed a *probatio diabolico*, a devil’s burden of proof, because it is very hard to prove). In the case of strict liability, such counter-proof is not allowed.

To decide whether Coca Cola was liable for the harm to Duchateau, because of Sjouwerman’s behaviour (vicarious liability), the court must decide whether the act of Sjouwerman was a violation of an unwritten ‘duty of care’. The court of first instance decided that Sjouwerman was not at fault, because Duchateau should have been more careful himself. The court of appeal, however, found Sjouwerman at fault, notably for not taking into account that customers may not be as prudent as might be expected. Considering the major consequences of an accident, Sjouwerman should have taken safety measures to prevent this.

The Supreme Court found that the court of appeal had used correct criteria to assess whether Sjouwerman violated his duty of care towards the customers of the cafe, notably:

1. the probability that visitors of the cafe are not as cautious as necessary; and
2. the probability that this lack of caution will lead to accidents; and
3. the seriousness of the harm that may result; and
4. the extent of the burden of safety measures.

Here we see that private law contains a number of generic concepts, such as ‘duty of care’ that require a case-specific assessment of what is at stake, while taking into account that the assessment criteria must be *generalizable to subsequent cases*—in line with both legal certainty and justice (treating similar cases similarly in a foreseeable manner). In Chapter 8 we will revisit tort law in more detail, in relation to privacy harms and cyber torts.

Legal judgment is a crucial but complex and reflective practice, demanding acuity and ingenuity of the court in the face of changing circumstances and the competing demands of legal certainty, instrumentality, and justice. It highlights (1) the need to assess and interpret the facts of the case in light of the applicable legal framework, and (2) the simultaneous need to identify and interpret the applicable legal norm in light of the legal framework and the facts of the case. All this

demonstrates the inherent contestability of judgments, both regarding the identification of relevant facts and the interpretation of the legal norm. In turn, all this highlights the centrality of both interpretation and legal reasoning in the study and the practice of law.

3.3 Public Law and Criminal Law

As discussed above, decisions under public law must be justifiable in terms of the public interest. It may be that it is in the public interest that the state considers and defends its own interests, for instance in situations of emergency. This, however, cannot be assumed: the interest of the state should not be conflated with the public interest.

In a *constitutional democracy*, the state must not only

- act with an eye to the *public interest*; but also
- act within the confines of the *legality principle*; and
- treat citizens with equal respect and concern.

These requirements similarly apply to criminal law, which involves one of the most invasive competences of the state, namely the so-called *ius puniendi* (the right to punish). Legal scholarship often qualifies criminal law as a subdomain of public law, as it constitutes and regulates the conduct of the state. It contains the secondary rules that clarify which primary rules are protected by means of criminalization. This can be gleaned from the articulation of criminal offences, for example: ‘Whoever commits murder will be punished with maximum 15 years of imprisonment.’ This is clearly not a primary rule; it does not state that murder is prohibited. Instead, it states under what legal conditions punishment is lawful. Criminal law, in that sense, depends on the vertical relationship between a state and its citizens—as in public law.

However, the secondary rule manifestly assumes the primary rule; one cannot be punished if one’s conduct is not unlawful. That is why some jurisdictions do not qualify criminal law as a subdomain of public law, emphasizing that the primary rules concern the horizontal relationships between legal subjects. One could say that criminal law shows the mutual dependencies between the horizontal and the vertical relationships of a legal system.

3.3.1 Public law

Public law regards, on the one hand, legal relations between a state (acting as such) and its citizens, and, on the other hand, the legal relationships between states. The first concerns constitutional and administrative law, the second concerns international public law. Constitutional law and international public law have many dependencies, in the first place because the constitution determines if, to what extent, and under what conditions international law overrules national law in case of a conflict between both. Second, international law may stipulate its own priority, for instance in the case of *ius cogens*, that is, law that applies without exception to all states (e.g. the prohibition of crimes against humanity and genocide).

International public law will be discussed extensively in the next chapter (Chapter 4). Here, we focus on constitutional and administrative law.

3.3.1.1 Constitutional law

Constitutional law *attributes competences*:

1. to legislate (Acts of Parliament) and to regulate (Regulatory Policies);
2. to act and decide based on its public authority (traffic management, environmental protection, decisions on tax or social security); and
3. to adjudicate (private law, criminal law, administrative law).

These competences are attributed to the legislator (e.g. parliament, municipality), to public authorities (cabinet ministers, supervisors, tax authorities, environment agencies), and to courts (defining their jurisdiction).

Constitutional law *restricts the competences it attributes* by requiring specific safeguards which constitute legal conditions that limit the exercise of the powers that have been allocated. This clearly shows the *constitutive and limitative* nature of the attribution of powers in a constitutional democracy. These limitations may concern procedural or substantial prerequisites, for example, making sure that privacy is not unnecessarily infringed, unjustified discrimination is prevented, and the freedom of speech is not violated.

3.3.1.2 Administrative law

Administrative law regulates the conduct of the government and other agencies with public authority, for example, in the domain of environmental law, student grants law, public health law, and tax law.

Based on the *legality principle*, administrative law requires that actions and decisions of public authorities have a legal basis.

This legal basis *constitutes their competence* to, for example, maintain roads, to take decisions about individual taxes or social security, and to impose policy rules on the industry regarding specified pollution (emission thresholds). Citizens addressed by the decisions have a duty to obey, and such decisions are often assumed to be lawful, even though their lawfulness may be contestable in an appeal procedure.

The same legality principle *limits the competences* of public authorities, by making them conditional upon statutory constraints and safeguards. Next to this, some jurisdictions have developed unwritten principles that have the force of law, thus regulating how public authorities can use their competences. Such principles are often divided into substantive and procedural principles, for instance: the principles of *trust and legitimate expectations*, *fair play*, *reasoned decisions*, and *proportionality* as well as *subsidiarity*. These principles have been recognized and developed by courts with jurisdiction concerning actions and decisions based on administrative law. Such jurisdiction provides citizens with legal remedies against public authorities.

Legal remedies form a crucial safeguard in the context of administrative law, as they give citizens the competence to appeal against decisions of public authorities in a court of law. Such an appeal may or may not suspend the duty to comply with the decision and should, for example, enable the testing of the validity or applicability of the legal basis, as well as the manner in which the administration has used its competences. Imagine that the tax authorities decide that one's income over 2017 is €120,000, imposing the correlated income tax of, for example, €67,000, whereas another interpretation of what constitutes one's actual income results in an income of €110,000. Without a system of legal remedies, one could maybe ask the tax authorities to revise the decision, but lack *the right* to present one's position to an independent court. Once that court has decided, other taxpayers have a more precise understanding of how one's income should be calculated. So, the system of legal remedies in administrative law contributes to legal certainty.

Constitutional and administrative law are core to the rule of law; they vouch for the integrity of a government's conduct versus its citizens. We should remind

ourselves that *limited government* cannot be taken for granted. The idea that government must be brought under the rule of law, facing countervailing powers to rein in its potentially unlimited rule, is a historical artefact that must be reinvented and sustained. We may want to add that there is no need to be either naïve or cynical about the practice of limited government; the checks and balances of the rule of law must be instituted, reinvented, sustained, and vigilantly defended. Neither taking them for granted, nor cynically denying their impact will do.

3.3.2 Criminal law

3.3.2.1 Substantive criminal law

Criminal law is usually divided into *substantive law* and procedural law. The first contains the primary rules (prohibitions) that delineate which actions qualify as criminal offences. In the Criminal Code, as discussed above, these primary rules are often hidden in secondary rules, for example, ‘Whoever hacks into a computing system without right, can be punished with maximum four years of imprisonment and/or a fine of €20,000’. Though this norm explicitly addresses the state, attributing the legal power to punish a person if specified conditions are satisfied, the norm indirectly addresses citizens by delineating a prohibited action as punishable.

Outside the Criminal Code, for example, in Acts of Parliament that legislate on traffic, environmental, or tax law, the primary norms are formulated separately. Here, separate secondary rules impose criminal or administrative sanctions, often situated at the end of the Act in a chapter on enforcement. Note that in many jurisdictions the administration does not have the legal power to criminalize unlawful behaviour, unless specifically authorized and conditioned in an Act of Parliament.

An example of a secondary rule that criminalizes the violation of a separately formulated primary rule, would be section 118(1)(a) of the UK Environmental Protection Act 1990, Chapter 43:³

- 118 (1) It is an offence for a person—
- (a) to do anything in contravention of section 108(1) above in relation to something which is, and which he knows or has reason to believe is, a genetically modified organism;

³ <http://www.legislation.gov.uk/ukpga/1990/43>.

This is the secondary norm by which the violation of a primary norm is criminalized. The relevant primary norm can be found in section 108(1) of the same Act:

- 108 (1) Subject to subsections (2) and (7) below, no person shall import or acquire, release or market any genetically modified organisms unless, before doing that act—
- (a) he has carried out an assessment of any risks there are (by reference to the nature of the organisms and the manner in which he intends to keep them after their importation or acquisition or, as the case may be, to release or market them) of damage to the environment being caused as a result of doing that act; and
 - (b) in such cases and circumstances as may be prescribed, he has given the Secretary of State such notice of his intention of doing that act and such information as may be prescribed.

Substantive criminal law thus determines (1) which conduct is punishable, (2) with what punishment. To be punishable, conduct must at least be unlawful, but—as indicated above, this is not a sufficient condition (see Figure 3.3). To be punishable the relevant conduct must be defined in a way that clarifies in a precise way when citizens become liable to punishment. Legal certainty requires that the scope of the offence must be transparent to those subject to the legal effect of criminalization. Note that the legal effect is not punishment, but *punishability*.

Legal certainty is further enhanced and protected by the criminal law legality principle, which is even more stringent than the generic legality principle of public law.

In criminal law, this principle is also called the *lex certa* principle that safeguards: (1) a reasonably precise formulation of criminal offences to prevent overinclusive criminalization; and (2) protection against retroactive application. The latter entails that actions (including omissions) can only be punishable if they were criminalized when performed.

By way of example, we will discuss a leading case of Dutch case law, under the heading of ‘Old style smart metering’, though it usually goes by the name of ‘the electricity judgment’.⁴

The facts of the case are quite simple: a dentist in The Hague repeatedly uses a knitting needle to halt the electricity meter, thus reducing his electricity bill.

⁴ Netherlands Supreme Court, 23 May 1921, NJ 1921/564.

He is charged with theft and sentenced to three months of imprisonment. The criminal offence of theft is defined in Article 310 of the Netherlands Criminal Code (NCrC):

Who takes away a good that belongs in whole or in part to another, with the intention to appropriate it unlawfully, will, as guilty of theft, be punished with imprisonment of at the most 4 years or a fine of the fourth category.

The legal effects that apply, if the legal conditions are satisfied, are: (1) that one is guilty of theft, and (2) therefore punishable by way of an imprisonment of a maximum of four years, or a fine of the fourth category.

This legal effect depends on the following legal conditions:

1. a person has taken away
2. a good
3. that belongs to another (in whole or in part)
4. with the intention to appropriate it
5. unlawfully.

In this case, the dentist appealed to the Supreme Court, claiming that electricity is not a good, because it is not tangible. The Advocate-General (a formally appointed adviser to the Court) agreed and pointed out that other intangible goods such as intellectual goods cannot be stolen.⁵ The Supreme Court, however, decided that the term good should be understood to encompass electricity. It gave the following reasons for qualifying electricity as 'a good' in the sense of Article 310 NCrC: it can be transferred, accumulated and kept in store; it has an economic value; it can be taken away and appropriated unlawfully (by using the knitting needle).

In other words, the court stipulated the following criteria to qualify something as 'a good' in the sense of Article 310 NCC:

1. an independent existence;
2. transferability;
3. economic value; and
4. appropriation.

⁵ Note that the advice of the Advocate General is not binding upon the Court.

This case was a seminal case, because it seemed to reason by way of analogy: if money is a good, the same goes for electricity; if taking away another's money without right is a criminal offence, the same goes for taking away another's electricity. More abstractly, one could argue that the court said that since stealing a tangible good is punishable, stealing an intangible good is also punishable. This could have many consequences for the theft of other intangibles, such as intellectual property rights or other types of information.

In substantive criminal law, reasoning by way of analogy is prohibited. The reason is that this could extend the scope and the reach of the criminal law beyond what those subject to its legal effect legitimately expect. This is why the Court went out of its way to clarify that its reasoning is not a matter of *analogy*, but of *extensive interpretation*. Instead of saying that since stealing a tangible good is punishable, stealing an intangible good is also punishable, the court said that the term good must be understood to include electricity, even if it does not include intangible goods. To justify such extensive interpretation, the court provided a set of reasons that clarify that this interpretation is reasonable and fits the system and the purpose of the relevant law.

With the advent of digital data, the question of what qualifies as 'a good' in Article 310 (theft) and Article 321 (embezzlement) NCrC has returned many times. In the Netherlands, for instance, in cases about embezzlement of money from a bank account,⁶ 'stealing' data from another's computing system,⁷ stealing money with a smartcard and password via an ATM,⁸ and 'stealing' bandwidth.⁹ In the latter case the court found that 'taking' bandwidth does not imply that others have less, suggesting this may nevertheless qualify as an offence under Article 138ab NCrC (unlawful access to an 'automated work').

The Court of Appeal in The Hague tested this option,¹⁰ but concluded that since a router is not an 'automated work' Article 138ab NCrC does not apply. This was based on Article 80sexies NCrC, which stated that:

An 'automated work' is to be understood as a device that is meant to store, process and transfer data electronically.

⁶ Netherlands Supreme Court, 11 May 1982, NJ 1982, 583, where the court decided that scriptural money is a good in the sense of Art. 321 NCrC, because of its function in societal intercourse.

⁷ Netherlands Supreme Court, 3 December 1996, ECLI:NL:HR:1996:ZD0584, where the court held that such data does not constitute 'a good' because there is no loss of possession on the side of the holder of the data. This led to the legislator taking over, enacting a new criminal offence under Art. 138ab(2) NCrC which criminalizes unauthorized access to a computing system (with additional punishment for copying of data).

⁸ Netherlands Supreme Court, 19 April 2005, ECLI:NL:HR:2005:AS9237, where the court decided that using a stolen smartcard and pin code may qualify as theft 'using a false key'.

⁹ Netherlands Court of First Instance Amsterdam, 11 September 2008, ECLI:NL:RBAMS:2008:BF0824.

¹⁰ Netherlands Court of Appeal The Hague, 9 March 2011, ECLI:NL:GHSGR:2011:BP7080.

The court of appeal decided that a router is not an ‘automated work’ because it does not store, process, and transfer data. This, in turn, was overturned by the Supreme Court,¹¹ clarifying that a router is part of a networked computing system that can store, process, and transfer data (thereby it is an ‘automated work’).

Moving deeper into the onlife world, two more key judgments offer an interpretation of ‘a good’ in the sense of Article 310 NCrC. In 2012, when asked whether virtual goods (‘owned’ in an online game environment) can be stolen, the Supreme Court decided that, indeed, depending on the circumstances data can be qualified as a ‘good’ in the sense of Article 310 NCrC.¹² Also in 2012, when asked whether SMS-messages and mobile phone minutes can be stolen, the Supreme Court confirmed that this is possible.¹³

We can now draw a nice timeline, specifying the legal conditions that must apply to qualify something as a ‘good’ in the sense of Article 310/326 NCrC:

1. Independent existence (Electricity Judgment 1921);
2. Transferability (Electricity Judgment 1921);
3. Economic value (Electricity Judgment 1921);
4. Appropriation (Electricity Judgment 1921);
5. Function in societal intercourse (Money in a bank account Judgment 1982);
6. Loss of possession after transfer (Stealing data judgment 1996).

These conditions have been applied in all subsequent judgments, highlighting that to qualify as a good it must be rivalrous (one person having more implies another person having less) as well as exclusivity (either the victim or the perpetrator has control over the good). These criteria will co-determine answers to new questions, such as whether stealing a pin code via a brain interface qualifies as theft in the sense of Article 310 NCrC, or as unlawful access to a computing system under Article 138ab NCrC.

A similar case has been decided already in 1995, where the Supreme Court, applying the above criteria decided that a pin code in the mind of a person is not a ‘good’ in sense of Article 317 NCrC (concerning extortion, blackmail).¹⁴ This was a first inkling that ‘loss of possession’ is a critical condition to qualify

¹¹ Netherlands Supreme Court, 26 March 2013, ECLI:NL:HR:2013:BY9718.

¹² Netherlands Supreme Court, 31 January 2012, ECLI:NL:HR:2012:BQ9251.

¹³ Netherlands Supreme Court, 31 January 2012, ECLI:NL:HR:2012:BQ6575.

¹⁴ Netherlands Supreme Court, 13 June 1995, ECLI:NL:HR:1995:ZD0064.

something as a good. So, what if we could actually remove a pin code from a person's brain?

Back to why it matters whether a certain conduct does or does not fall within the scope of a criminal offence *at the time of the conduct*. Why prohibit retro-active criminalization and analogous reasoning, and why should extensive interpretation be limited? Because the legal effect of criminalization means that conduct becomes punishable, the invasive nature of punishment requires an enhanced degree of legal certainty (as discussed above this is called *lex certa* and aligns with the criminal law legality principle). At the end of the eighteenth century, the famous legal scholar Beccaria formulated this principle as a maxim for a legitimate criminal law: *nullum crimen, nulla poena sine preavia lege poenali* (no crime, no punishment without prior criminalization).

One final example is one that may speak to a computer scientist. In a tweet, Ted Neward wrote:¹⁵ 'Every. Single. Software developer. Must. Take. Note. YOU can go to jail for the code YOUR BOSS tells you to write.' He linked a news item about a Volkswagen engineer, who 'helped develop the software that concealed high levels of pollutants generated by Volkswagen's diesel engines'¹⁶ and who was subsequently sentenced to forty months' imprisonment. The sentence has been considered as harsh, because the engineer was not the mastermind of the deceptive scheme and merely seemed to have followed orders. The objective of the conviction was not only *retribution* (punished based on desert), but clearly also *deterrence* (punishment meant to warn off other engineers from following orders to commit a criminal offence). This raises many fascinating questions about the goals of legitimate public punishment, including the question of whether one offender may be used as an example to deter others and how this relates to justice and equal treatment if similar offenders are not prosecuted. Let us take note that it is not possible to prosecute each and every suspect of a criminal offence, while also raising a flag about the legitimacy of policies developed to make the right kind of choices in the course of public prosecution.

The strict requirements surrounding the articulation of a criminal offence also concern potential *justification* and *excuse*. Even when a person commits the offence as defined, they may be able to justify their action. For instance, one may have killed another person intentionally, thus fulfilling the legal

¹⁵ <https://twitter.com/tedneward/status/901135785969074177>.

¹⁶ Vlasic, Bill. 2017. 'Volkswagen Engineer Gets Prison in Diesel Cheating Case.' *The New York Times*, 22 December 2017, available at: <https://www.nytimes.com/2017/08/25/business/volkswagen-engineer-prison-diesel-cheating.html>.

conditions of manslaughter, but nevertheless not be punishable because: (1) the act was justified by *self-defence*, having to prevent one's own or another's death or serious injury, or (2) the act must be excused because the situation caused such *overwhelming psychological stress* that one cannot be considered guilty for having killed a thief who nevertheless did not threaten one's or another's life. Though justification and excuse cannot be taken lightly and will require serious argumentation, to be convicted for a criminal offence, both *wrongfulness of the act* and *culpability of the perpetrator* must be confirmed.

This leads to a quadruple structure of a criminal offence, in other words, the legal effect of punishability is conditional upon the following legal conditions:

Actus reus (the act and its qualification):

1. an *action* (in criminal procedure this relates to the law of evidence),
2. that falls within the scope of a *criminal offence* (in criminal procedure this regards the qualification of conduct as a criminal offence).

Mens rea (the elements):

3. *wrongfulness* (in criminal procedure this regards the defence of justification);
4. *culpability* (in criminal procedure this regards the defence of disculpation or excuse).

3.3.2.2 Criminal procedure, including police investigation

As indicated, the structure of the criminal offence is deeply entwined with criminal procedure, notably with the questions a court must answer before convicting a defendant. These questions highlight the crucial role played by contestability at the heart of the law.

The legal effect of a conviction thus depends on all of the following questions being answered positively (the conditions are cumulative).

1. The conduct that is charged must be *proven* beyond reasonable doubt.
Defence: 'I did not do it.'
2. The proven conduct must *qualify as a criminal offence* (legality principle).
Defence: 'The proven conduct is not a criminal offence.'
3. The action was *wrongful* (no justification).
Defence: 'I had a ground for justification' (e.g. permission).

4. The defendant was *culpable* (no excuse).
 Defence: 'I had an excuse' (e.g. psychiatric disorder).

In the case law about whether data, bandwidth, or virtual goods qualify as 'a good' in the sense of Article 310 NCrC, the second question was at stake.

Criminal procedure concerns both pre-trial police investigations and the trial itself. The legality principle that informs government competences under public law also applies to the police, public prosecutor, and the courts when deciding criminal cases. Due to the impact of punishment and the invasive character of criminal investigation, legality issues in criminal procedure are core to the legitimacy of criminal investigation, prosecution, and conviction.

In the context of criminal procedure, the term 'legality principle' is also used to denote a strict form of legality, referring to the idea that all criminal conduct should be prosecuted. This interpretation of the legality principle is opposed to the idea that the public prosecutor has *discretion* when deciding whether or not to prosecute. In the Netherlands, for instance, this discretion is codified in the Code of Criminal Procedure (NCCrP), stating that the prosecutor may abstain from prosecution based on the general interest.¹⁷ Case law has clarified that the office of the public prosecutor must develop a policy, specifying the criteria that are used to determine whether or not to prosecute. Examples can be found in the Dutch policies around euthanasia and the possession of soft drugs. In both cases, the relevant actions (of a doctor performing euthanasia or a person walking around with soft drugs) remain criminal offences, but the office of the public prosecutor has developed and published *policy rules* that detail under what conditions a doctor or a soft drug user will not be prosecuted. *Legal certainty* not only demands that citizens can foresee which of their conduct is punishable, but can also foresee under what conditions they will be prosecuted.

The difference between:

1. a strict legality principle that requires prosecution of all alleged criminal offences, and
2. a principle of discretion that makes room for policy considerations, connects with different justifications of punishment.

¹⁷ Article 167.2 NCCrP.

Some theories highlight that punishment is *retribution* for the violation of norms that must be upheld in the general interest, even if no concrete, identifiable damage has been caused. This would rule out any discretion to abstain from prosecution. Other theories highlight that punishment is meant to *prevent* further crime, both by way of deterring others from committing similar offences (general prevention) and by way of preventing the convicted offender from re-offending (specific prevention). Most jurisdictions are based on a combination of retribution and prevention; the public prosecutor must develop and publish its policies to clarify how *discretion* will be exercised. Without such policy the decision to prosecute could be arbitrary, depending on private considerations of whoever holds the office of the public prosecutor instead of justifiable choices with regard to the public interest. Note that in practice it is not even remotely possible to prosecute all criminal offences. Acknowledging this and being transparent about the foreseen use of discretionary competences form important legal safeguards against *arbitrary* punishment.

In Europe, a criminal charge results in the applicability of the right to a fair trial, as articulated in Article 6 of the European Convention on Human Rights (ECHR):

1. In the determination of his civil rights and obligations or of any criminal charge against him, everyone is entitled to a fair and public hearing within a reasonable time by an independent and impartial tribunal established by law. Judgment shall be pronounced publicly but (...).
2. Everyone charged with a criminal offence shall be presumed innocent until proved guilty according to law.
3. Everyone charged with a criminal offence has the following minimum rights:
 - (a) to be informed promptly, in a language which he understands and in detail, of the nature and cause of the accusation against him;
 - (b) to have adequate time and facilities for the preparation of his defence;
 - (c) to defend himself in person or through legal assistance of his own choosing or, if he has not sufficient means to pay for legal assistance, to be given it free when the interests of justice so require;
 - (d) to examine or have examined witnesses against him and to obtain the attendance and examination of witnesses on his behalf under the same conditions as witnesses against him;
 - (e) to have the free assistance of an interpreter if he cannot understand or speak the language used in court.

We can detect six fair trial principles that underlie this right:

1. the presumption of innocence;
2. the right to an independent and impartial tribunal;

3. equality of arms between public prosecutor and defendant, including internal publicity;
4. immediacy of the presentation and testing of the evidence in court;
5. external publicity; and
6. the right to have a final decision within a reasonable time.

Taken together, these rights ensure that a defendant has the means to *contest* the lawfulness of police investigations and the evidence presented by the prosecution, including witness testimony. They also make sure that in principle the *burden of proof* is on the public prosecutor and until guilt has been established the defendant is not to be treated as if he is a perpetrator. This means that all measures taken before a conviction must serve other purposes than punishment; they should not be deterrent or punitive. Together these requirements condition lawful investigations and prosecution, and a valid conviction. They have generated a steady flow of case law of the European Court of Human Rights (ECtHR) that has jurisdiction to hear individual complaints of citizens of the Contracting Parties of the Council of Europe that instigated the ECHR.

This case law has, for instance, determined that the term ‘criminal charge’ has an autonomous meaning which does not depend on what a state defines as punitive sanctions. States therefore cannot disable the applicability of Article 6 ECHR, e.g. by re-naming criminal offences as ‘regulatory offences’.

If they were to label criminal offences as ‘regulatory offences’, Article 6 ECHR nevertheless applies if:

- the nature of the offence, and
- the severity of the penalty,

bring the offence within the bounds of the concept of ‘a criminal charge’.

This will depend, for example, on whether the sanctions have punitive and/or deterrent objectives, or on whether other Contracting Parties qualify the offence as a criminal offence.¹⁸ Other case law determined that defendants must have access to and be able to challenge all and any evidence presented to the court, even if the public prosecutor wishes to hide information on grounds of

¹⁸ ECtHR, 8 June 1976, Application no. 5100/71; 5101/71; 5102/71; 5354/72; 5370/72 (*Case of Engel and Others v. The Netherlands*).

public interest.¹⁹ More recently, the ECtHR decided that suspects that are interrogated by the police have a right to legal counsel.²⁰

Though we did not discuss private law procedure it makes sense to say a few words at this point about private law procedure as compared to the criminal trial.

1. First, let's take note that the right to a fair trial (Article 6(1) ECHR) also applies to the determination of a person's civil rights and obligations, whereas Article 6(2) and (3) are reserved for a criminal charge.
2. Second, in private law proceedings, the default rule is that whoever initiates proceedings bears the *burden of proof*. Think of requesting a court order to comply with contractual obligations, an injunction to stop unlawful conduct, or compensation for damage caused by a breach of contract or a tort. In case of liability for high-risk conduct, the burden of proof is sometimes inversed, while risk liability and strict liability may further diminish the burden for the plaintiff. Think of the use of asbestos or other pollutants by the industry, which have been proven to cause grave health problems, or safety hazards in employment situations. Legislatures and courts thus aim to provide effective protection for victims, especially when causality can be inferred at a statistical level (increased probability to suffer harm or damage), but not determined at the individual level (where, e.g. other causes may have contributed to the damage). In the criminal trial, the public prosecutor bears the burden of proof, as part of the presumption of innocence (Article 6(2)).
3. Third, whereas the presumption of innocence demands that in a criminal procedure *the standard of proof* is 'beyond reasonable doubt', in private law the standard is usually much lower, for example, clear-and-convincing evidence or even preponderance of evidence. Plausibility is often considered enough.
4. Fourth, if the defendant in private law proceedings does not contest the evidence, the plaintiff's request must normally be granted. This goes back to the idea that within private law, parties are treated as autonomous and equal persons, capable of deciding amongst themselves the scope and the shape of the conflict. Such *party autonomy* does not exist in the criminal law, where imposing punishment on an innocent person is to be avoided even if defendant and prosecutor were to strike a deal.²¹ Since criminal law attributes the state

¹⁹ ECtHR, 16 February 2000, Application no. 28901/95 (*Case of Rowe and Davis v. the United Kingdom*).

²⁰ ECtHR, 27 November 2008, Application no. 36391/02 (*Case of Salduz v. Turkey*).

²¹ This is clearly different in e.g. the United States, where such deals are a regular way of managing the case load of the courts. Even in continental European legal systems public prosecutors may have far reaching competences to reach an agreement with a defendant; if the case goes to court, however, the court must establish the facts—irrespective of deals struck by the prosecutor.

with a number of invasive legal powers, a more active position of the court is warranted when it comes to deciding the reliability and the relevance of the evidence and its contribution to proving the offence ‘beyond reasonable doubt’. In the criminal trial the defendant and public prosecutor are not considered equal, calling for a set of compensatory rights to provide the defendant with effective means to defend themselves.

References

Private law

For a general but more detailed introduction to private law

Hage, Jaap, Antonia Waltermann, and Bram Akkermans, eds. 2017. *Introduction to Law*. 2nd ed. New York: Springer, chapters 3 (conceptual distinctions), 4 (contract law), 5 (property law), and 6 (tort law).

On the legal order as a system of legal relationships and the concept of individual rights

Achterberg, Norbert. 1982. *Die Rechtsordnung als Rechtsverhältnisordnung: Grundlegung der Rechtsverhältnistheorie*. Berlin: Duncker & Humblot.

Edmondson, William A. 2012. *An Introduction to Rights*. 2nd ed. Cambridge: Cambridge University Press.

For a comparative perspective on tort law, notably the unwritten duty of care

Van Maanen, Gerrit, David Townend, and Almaz Teffera. 2008. ‘The Dutch “Cellar Hatch” Judgment as a Landmark Case for Tort Law in Europe: A Brief Comparison with English, French and German Law with a “Law and Economics Flavour”’. *European Review of Private Law* 16 (5): 871–89.

On common principles and rules of European private law

Bar, Christian von, and Eric Clive, eds. 2010. *Principles, Definitions and Model Rules of European Private Law: Draft Common Frame of Reference*. Oxford and New York: Oxford University Press.

More specifically European comparative property law

Erp, Sjef Van. 2006. ‘Comparative Property Law’. *The Oxford Handbook of Comparative Law*, November. <https://doi.org/10.1093/oxfordhb/9780199296064.013.0033>.

More specifically European contract law

Kötz, Hein. 2017. *European Contract Law*. 2nd ed. Oxford and New York: Oxford University Press.

On fundamental principles of private law from a constitutional theory perspective

Mak, Chantal. 2012. 'Europe-Building Through Private Law: Lessons from Constitutional Theory'. SSRN Scholarly Paper ID 2023141. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=2023141>.

For an in-depth understanding of the freedom to act strategically in the private sphere

Habermas, Jürgen. 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Studies in Contemporary German Social Thought. Cambridge, MA: MIT Press.

Public law

A general but detailed introduction to constitutional, administrative, and international public law

Hage, Jaap, Antonia Waltermann, and Bram Akkermans, eds. 2017. *Introduction to Law*. 2nd ed. New York: Springer, chapters 8 (constitutional law), 9 (administrative law), and 12 (international law).

On constitutional law

Frankenberg, Günter. 2012. 'Comparative constitutional law'. In *The Cambridge Companion to Comparative Law*, edited by Mauro Bussani and Ugo Mattei, 171–90. Cambridge: Cambridge University Press.

Waluchow, Wil. 2018. 'Constitutionalism'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2018/entries/constitutionalism/>.

On administrative law

Bignami, Francesca. 2012. 'Comparative Administrative Law'. In *The Cambridge Companion to Comparative Law*, edited by Mauro Bussani and Ugo Mattei, 145–70. Cambridge: Cambridge University Press.

Stroink, F., and E. Van der Linden, eds. 2005. *Judicial Lawmaking and Administrative Law*. Antwerpen and Maastricht: Intersentia.

On criminal law

A general but detailed introduction to criminal law

Hage, Jaap, Antonia Waltermann, and Bram Akkermans, eds. 2017. *Introduction to Law*. 2nd ed. New York: Springer, chapter 7.

Horder, Jeremy. 2016. *Ashworth's Principles of Criminal Law*. 8th ed. Oxford: Oxford University Press.

A discussion of the legality principle in the context of European criminal law

Peristeridou, Christina. 2015. *The Principle of Legality in European Criminal Law*. Cambridge: Intersentia.

Guidance on the right to a fair trial

European Court of Human Rights (ECHR). 2014. Guide on Article 6 of the European Convention on Human Rights. Right to a fair trial (criminal limb). <https://www.echr.coe.int/Pages/home.aspx?p=caselaw/analysis/guides&c=#>.

4

International and Supranational Law

There was a time when international law was considered a minor and separate subject in the study of law.

By the beginning of this century it was clear that merely studying one's own national law was not merely 'provincial' but also meant not being up to standards regarding positive law. The reason was that positive law, that is, valid and applicable law here and now, depends on national jurisdiction and, at least within Europe, national jurisdiction increasingly incorporates both international and supranational law.

For instance, fundamental rights are not only part of the national constitution, but can also be invoked based on the European Convention on Human Rights (ECHR) and—since 2009—based on the Charter of Fundamental Rights of the European Union (CFREU). Next to these human rights instruments many treaties have been concluded under international law on other subjects (e.g. the Cybercrime Convention (CC), the Trade-Related Aspects of Intellectual Property Rights (TRIPs) Agreement, the Convention on the Recognition and Enforcement of Foreign Judgments in Civil and Commercial Matters, and an entire body of supranational law (i.e. the law of the European Union (EU), such as the Copyright Directive, the Unfair Commercial Practices Directive, the Machinery Directive, which has become part of national jurisdiction in the Member States of the EU.

Clearly, the relevance of law for computer scientists—the architects of our new onlife world—cannot be reduced to that of one national jurisdiction. The combination of networked computational systems and the hyperconnectivity of the current information and communication infrastructure call for a keen acuity with regard to national, international, and supranational law.

In this book we will focus on international law in the context of the Council of Europe (CoE, forty-seven contracting states) and on supranational law in the context of the EU (twenty-seven Member States). As this book aims to provide insight in 'how lawyers think' and 'what law does,' not a comprehensive overview (which would be entirely undoable), we restrict ourselves to the most relevant legal instruments within the jurisdictions of Europe, as a good point

of entry. I believe this does not make the book less interesting for, for example, US, Australian, or even Asian computer scientists. On the contrary, this book aims to provide a coherent framework for *understanding how law operates*, combining analytical rigour and interpretive salience with concrete examples to demonstrate the relevance of the distinctions made and the perspectives taken. It would be great to add other jurisdictions as new examples, enriching the conversation at the global level on how to order our interactive, dynamic, and potentially turbulent world. Besides that being impossible in the context of one book, we must note that this book also takes a normative perspective on how law ought to operate, as part of a constitutional democracy, rejecting both instrumentalist or moralistic conceptions of law (see above, section 2.2.2). Even though European law is not in any way perfect, the European attempt to institute, sustain, and reinvent a moderate government that respects human rights is a good example of what law and the rule of law ‘do’.

In this chapter, we first discuss the concept of jurisdiction and its formative status in national, international, and supranational law, after which we provide a more in-depth overview of international law and supranational law.

4.1 Jurisdiction in Western Legal Systems

The concept of jurisdiction first appears in the early fourteenth century, and though it has tied in with the concept of territory, the latter term first appeared in the early fifteenth century. Even if Western legal systems equate jurisdiction with territorial jurisdiction this is not necessarily correct.

Actually, the concept of jurisdiction is often used in two different ways, as either:

1. the competence to legislate, adjudicate, and enforce; or
2. the territory or domain over which an entity holds jurisdiction in the first sense.

Both are relevant, and we can add a second distinction, with regard to:

1. internal jurisdiction, that is, the competence to legislate, adjudicate, and enforce the law within the state;
2. extraterritorial jurisdiction, that is, the competence of one state to legislate, adjudicate, or enforce its law on the territory of another state.

In Anglo-American discourse the term ‘power’ is used where Europeans use ‘competence’. Very simply defined, we could say that jurisdiction refers to legal power and to where such power is applicable. This raises challenging questions, such as to what extent a state can decide the limits of international or supranational jurisdiction on its own territory, and to what extent an international court gets to decide this. In other words: where must we situate the competence to decide the attribution, content, and limits of competence? Because German scholarship has worked on this, we call this the question of *Kompetenz-Kompetenz*.

4.1.1 An example

To sensitize the reader to issues of *Kompetenz-Kompetenz* I will take them through some of the issues encountered in international private law, which is in point of fact national law. What happens if Alies (Dutch) marries Bob (a US citizen) in Japan, but they will live in Russia? What law applies to the marriage: Dutch, US, Japanese, or Russian law? If they want to get divorced, which court is competent: a Dutch, a US, a Japanese, or a Russian court? What if they want their Dutch divorce to be recognized in Iran?

International private law confronts three types of questions:

1. the applicable law;
2. the competent court; and
3. enforcement.

The questions regarding *applicable law* ask which national law determines the legal consequences of the marriage. This may depend on choice, on a treaty, and in the end, it will always depend on national law, as the national law must recognize the choice (which may be guaranteed in a treaty signed by the relevant state). Note that the *primacy of national law* implies that a person may be married according to the national jurisdiction of the country where she lives, even after obtaining a valid divorce in another national jurisdiction.

The second type of questions concerns jurisdiction in the sense of *adjudicatory competence*. If one gets married in Russia, under Japanese law, which court is competent to decide on a divorce? Does this depend on the applicable law, on one’s residence, nationality, or on the country where the marriage took place?

The third type of questions concerns *recognition and enforcement*, asking under what conditions a court's divorce decision will be recognized and enforced in another country.

All these questions apply to issues of family law, as in the given example, but also to international sale of goods or services, capital investment, to labour conditions in transnational companies, or to keeping bank accounts in various countries. The complexity of the potential answers to these questions highlights the necessity of international treaties to reduce the uncertainty that evolves from this complexity. This regards questions of family law, property law, contract law, and tort law, and the global economy would be substantially disrupted without international treaties that bind the contracting parties (states), thus achieving a higher level of trust and legitimate expectations between citizens, companies, and other institutions that interact at the transnational level.

In the case of the marriage, one could wonder whether all this matters, or why we should care. Since a valid marriage has legal effects the answers to questions of international private law make a substantial difference. In some jurisdictions the default is that one marries on equal terms, which means that creditors of one partner have a legal remedy against the assets of the other partner. In other jurisdictions the default is that one marries under a separate estate arrangement, meaning that creditors of one partner have no legal remedy against the assets of the other partner. These defaults, as well as the possibility to opt for one or the other marital regime, differ in alternative national legal systems, and the same goes for the requirements for overruling the default regime (such as involving a notary public and the registration of prenuptial agreements).

4.1.2 National jurisdiction

What if the Netherlands want to delete the first article of their Constitution? What if the Netherlands wish to protect their citizens against internet activities undertaken in Russia or the United States by means of remote hacking by Dutch police officers? What if the Netherlands wish to abide by an overall *minimum* term of imprisonment of not more than one day and by Article 9a of the Netherlands Criminal Code (NCC);¹ can the

¹ Article 9a NCC reads: 'The court may determine in the judgment that no punishment or measure shall be imposed, where it deems this advisable, by reason of the lack of gravity of the offence, the character of the offender, or the circumstances attendant upon the commission of the offence or thereafter.'

Netherlands resist, for example, EU legislation that imposes higher minimum sanctions?

This section will discuss the *primacy as well as the limits of national jurisdiction* and its relationship to international and supranational law. To ensure legal certainty, lawyers need priority rules to determine the validity of legal norms whenever they are incompatible. The simplest way to achieve this is to assume that law is a hierarchical system of legal rules, where higher rules overrule lower rules. For instance, rules derived from the Constitution will overrule rules derived from Acts of Parliament, which in turn overrule rules derived from other public authorities with rule-making competences (e.g. municipalities, supervisors). This hierarchy works relatively well within the context of a single state. The reason is that each state has both internal and external sovereignty. The concept of sovereignty in this particular sense, stems from the 1648 treaty that introduced the so-called *Peace of Westphalia*. This treaty brought an end to a long and devastating period of European wars during the sixteenth and seventeenth centuries. These wars were both intra- and interstate and were entangled with *religious wars* between Roman Catholic and Protestant rulers, aiming to consolidate their own power over their subjects, based on adherence to their own religious allegiance. The Peace of Westphalia basically declared religion a matter of private faith and private consent, establishing the idea of a nation state with consolidated borders, where the sovereign holds the power to legislate, govern, and adjudicate within their territory (*internal sovereignty*), while respecting all other sovereigns as exclusively competent within their territory (*external sovereignty or the principle of non-interference*). Note that ‘the sovereign’ is not a person, but an office. It is this *office* that is competent, not the person that takes office. This institutionalization of sovereignty as an abstract entity that rules over an abstract geographical space still forms the root of the current system of sovereign states.

Both the sovereign and the territory are abstractions as they no longer depend on whoever takes the office of sovereign or whoever actually lives within the territory.

From 1648, one could say, the nation state takes centre stage, grounded by the idea of internal and external sovereignty—which form two sides of the same coin: without external sovereignty the sovereign cannot hold on to their internal sovereignty; without internal sovereignty the sovereign cannot ensure external sovereignty. The result is that *international law* becomes the law between independent sovereign states and thus depends on *consensus* between

these states. This is where supranational law fundamentally differs from international law, as supranational law depends on a partial transfer of sovereignty (conferral).

One of the assumptions of the current system of sovereign states is

- (1) that states can only be bound by international or supranational law if they so decide, as sovereigns can only obey rules outside their jurisdiction if they have bound themselves to those rules.

However,

- (2) the powerplay between states and between states and other powerful players, such as transnational companies and organizations, challenges the assumption of self-sovereign statehood.

Moreover,

- (3) various rules of international law do not depend on consent of individual states, but on assumptions about what constitutes lawful conduct, irrespective of sovereign will (*ius cogens*, fundamental principles of international law and some instances of customary international law).

In the case of *supranational law*, things become even more complex, because contracting states *give up part of their sovereignty* to enable effective collaboration and coordination within the jurisdiction of the EU.

It is therefore also crucial to remember that:

- (4) whereas the national jurisdictions of individual states are mutually exclusive, national, international and supranational jurisdiction will often overlap, and
- (5) the sovereignty of states depends on a system of international law that both *assumes and attributes* such sovereignty (see section 4.4 below).

4.2 International Law

It should be clear from the previous section that the actors in the domain of international law are, first of all, sovereign states. However, by now, other actors are recognized as such: international organizations (e.g. World Trade Organization (WTO), the UN), multinational companies (e.g. Shell, Google),

non-governmental organizations (NGOs) (e.g. Greenpeace), and even individuals as bearers of rights under international law.

4.2.1 Sources of international law

In international law, as in domestic (national) law, the sources of law determine the identification of the applicable legal norms. Because—in principle—international law is dependent on the consent of sovereign states, treaties are an obvious source of international law. Examples of international treaties are the Cybercrime Convention (CoE) 2001, the Berne Convention (copyright) 1971, the Paris Convention (patents etc.) 1883, the TRIPs Agreement (WTO) 1994, the International Covenant on Civil and Political Rights (UN) 1966, and the ECHR (CoE) 1950.

Treaties, however, are not the only source of international law. Article 38 of the Statute of the International Court of Justice in The Hague, states the following:

1. The Court, whose function is to decide in accordance with international law such disputes as are submitted to it, shall apply:
 1. international conventions, whether general or particular, establishing rules expressly recognized by the contesting States;
 2. international custom, as evidence of a general practice accepted as law;
 3. the general principles of law recognized by civilized nations;
 4. subject to the provisions of Article 59, judicial decisions and the teachings of the most highly qualified publicists of the various nations, as subsidiary means for the determination of rules of law.

The International Court of Justice (ICJ) was established by the UN Charter, which was signed immediately after the Second World War, in 1945. It is composed of fifteen judges and settles legal disputes between states and gives advisory opinions on legal issues to organs and agencies of the UN. Only states may appear before and apply to this Court and the Court can only settle disputes if both parties have recognized its jurisdiction by way of a declaration 'that they recognize as compulsory *ipso facto* and without special agreement, in relation to any other State accepting the same obligation, the jurisdiction of the Court' (Article 36, para. 2, of the Statute of the ICJ). Most textbooks on international law will summarize the sources of international law as:

- Customary law (*usus, opinio necessitatis*)
This regards not just any 'habit' or 'regularity in behaviour' but a combination of a particular state practice (*usus*) and the recognition that such a practice expresses a legal obligation (*opinio necessitates*).
- Treaties
This regards 'contracts' between states, based on the end result of negotiated text, *signed* by the representatives who negotiated the text and *ratified* by the heads of state, after internal agreement within the states. Normally treaties enter into force after a set number of ratifications.
- General principles of law
This regards, for instance, promotion of human rights and *self-determination of a people*, strict limitation of the use of force against other states, strict prohibition of acquisition of territory of another state by means of force, *principle of non-intervention*, and equality of states.
- Judgments and doctrine
This regards judgments of international tribunals and doctrine as published by respected scholars in international law.
- Decisions of international bodies
This regards decisions of, for example, the WTO or specialized bodies of the UN.
- Unilateral actions or declarations of states
Insofar as it is based on consensus, international law must accept state practice that, for example, rejects specific claims of customary law, and accepts declarations by states that reject the implications of judgments by Courts whose jurisdiction they do not accept.
- *Ius cogens*, obligations *erga omnes*
These are considered independent of the consent of states, as they concern the most flagrant violations of human dignity, genocide, and crimes against humanity. This implies that even unilateral actions or declarations of individual states cannot absolve them from the applicability of *ius cogens* (peremptory legal norms). Obligations *erga omnes* means that these obligations are absolute (for every state, regarding every other state or person).

4.2.2 Monism and dualism in international law

How does international law bind a state that is subject to its jurisdiction? And under what conditions does international law have direct effect, that is, direct legal effect for citizens in the form of providing them with legal rights?

Legal doctrine makes an analytical distinction between two approaches to the relationship between national and international law: a monist approach and a dualist approach.

A monist approach recognizes only one hierarchical legal order, of which international and national law form two parts and where international law has precedence over national law. As a consequence, in this approach, international treaties overrule national law and they have binding force as they are ratified, while citizens can appeal directly to international law, which national courts are legally bound to apply.

A dualist approach denies that national and international law are part of the same jurisdiction; they are considered as separate legal orders. To gain binding force within the national legal order, international law must first be transposed into national legislation. In this approach, citizens cannot directly appeal to international law but have to wait for its transposition, while the same goes for national courts, which are then only bound by national law.

The distinction is analytical and helps to understand the messy reality of overlapping national and international jurisdictions from the perspective of national law, which ultimately decides on the force of international law within its jurisdiction. In practice these approaches are both ends of a spectrum, with for instance the United Kingdom taking a dualist perspective and the Netherlands taking a mitigated monist (or a mitigated dualist) perspective.

The choice for a monist/dualist and mitigated perspective has far-reaching implications, which can be best understood in terms of legal effect. For instance, if we ask about the *legal effect of a treaty* that has come into force but has not been transposed into national law, the answer is that under a monist legal system national courts will have to apply the treaty, and the state may become liable *to its citizens* to the extent that it does not comply with the treaty. It thus has direct effect in the national legal order. Under the dualist legal system, the answer would be that national courts can only apply national law, and the state will become liable to the other contracting parties for non-compliance. The treaty will not have any direct effect in the national legal order.

As an example, let's check the Netherlands Constitution, Article 93:

Provisions of treaties and of resolutions by international institutions, which may be binding on all persons by virtue of their contents shall become binding after they have been published.

The phrasing of ‘binding on all persons by virtue of the content’ is equivalent with the concept of ‘direct effect’. The Netherlands Constitution basically states that any legal norm (of international law) which directly addresses legal subjects (corporations, natural persons) has legal effect for those legal subjects, who can invoke that norm in a national court of law. For legal norms with ‘direct effect’, the Netherlands implements a monist approach. Such ‘direct effect’, however, does not apply when a legal norm of international law addresses the contracting states instead of their citizens, thus imposing an obligation on states to enact the norm. In that case, the Netherlands employs a dualist approach. Article 93 thus follows the intent expressed in a treaty, identifying whether or not the treaty intends to directly create rights for citizens of the contracting parties.

Article 94 of the Netherlands Constitution clarifies even more clearly the hierarchical implications of its monist approach (in case of ‘direct effect’):

Statutory regulations in force within the Kingdom shall not be applicable if such application is in conflict with provisions of treaties that are binding on all persons or of resolutions by international institutions.

A prime example of a treaty with ‘direct effect’ is the ECHR.² Article 94 clearly shows that the ECHR must be applied by Dutch Courts, even if that results in the inapplicability of national law. This has wide-ranging consequences for the competence of Parliament, whose Acts can thus be overruled to the extent that they conflict with the human rights treaty. This is especially interesting due to the prohibition to test Acts of Parliament against the Constitution itself, as stipulated in Article 120 of the Netherlands Constitution:

The constitutionality of Acts of Parliament and treaties shall not be reviewed by the courts.

In the end, national legislation may be tested against provisions in international treaties with direct effect—but not against the Constitution. As one can imagine, this prohibition has been controversial and many attempts have been made to remove it from the Constitution. The argument in favour of this prohibition is that it clarifies the prerogative of the democratic

² Council of Europe, European Convention for the Protection of Human Rights and Fundamental Freedoms, as amended by Protocols Nos. 11 and 14, 4 November 1950, ETS 5, available at: <http://www.refworld.org/docid/3ae6b3b04.html> [accessed 13 October 2018].

legislature who should be the ultimate judge of whether an Act violates the Constitution.

Obviously, the Netherlands cannot be bound by international treaties unless its democratic legislature has consented. After a treaty is agreed and signed by the contracting parties, Parliament will have to decide whether or not the Netherlands will be bound by it. If Parliament consents, the head of state (the King) will ratify, binding the state to the treaty once it comes into force. This is worded in Article 91 of the Netherland Constitution:

The Kingdom shall not be bound by treaties, nor shall such treaties be denounced without the prior approval of the Parliament. The cases in which approval is not required shall be specified by Act of Parliament.

The manner in which approval shall be granted shall be laid down by Act of Parliament, which may provide for the possibility of tacit approval.

Any provisions of a treaty that conflict with the Constitution or which lead to conflicts with it may be approved by the Chambers of the Parliament only if at least two-thirds of the votes cast are in favour.

A prime example of a treaty without direct effect is the Cybercrime Convention (CC), which addresses the contracting states of the CoE and the other signatories, obliging them to enact a number of criminal offences and criminal law investigative measures in order to harmonize the criminal law enforcement measures against cybercrime. Neither the police nor individual defendants can invoke the CC directly, both will have to rely on the national implementation of its content by way of, for example, amendments of their Criminal Code and their Code of Criminal Procedure.

An important question with regard to the application of treaties, whether they have direct effect or require national implementation, is their interpretation and who gets to decide it: an international court, national court, or—to make things more complicated—both. If a treaty is concluded within a specific international jurisdiction, national courts may be bound to interpret the treaty in alignment with the case law of the relevant international court or tribunal. We can, for example, think of the ECtHR (relevant for contracting parties of the CoE) or of the ICJ (relevant for contracting parties of the UN). The interpretation of treaties will often involve the use of *preparatory documentation* that clarifies the intentions of the contracting parties and the underlying goals the treaty aims to support. An important source of law is constituted by the

preamble of a treaty, consisting of the so-called ‘recitals’ that articulate shared assumptions, goals, and explanations concerning the treaty. The *articles* of the treaty are considered binding law, they have legal effect (either direct effect for citizens of contracting parties, or direct effect for the contracting states). Such binding effect is missing for the *recitals*, but they are nevertheless an important source of law, as they provide authoritative information about how the articles should be read. Since international treaties are often the result of compromise, articles may be formulated in less clear terms, as this is often the only way to obtain agreement from all parties. The more radical text of previous drafts is sometimes moved to the recitals, thus leaving it up to the courts to decide the meaning of the article.

4.3 Supranational Law

Supranational law differs from international law. In the case of supranational law, a set of Member States (MSs) have agreed to transfer parts of their sovereignty to a supranational organization. In practice, supranational law refers to the law of the EU. Supranational law is not merely law between MSs (as in international law) but also law between the bodies of the EU and the citizens of the MS, who are also EU citizens. Some of the legal instruments of the EU have ‘direct effect’ for EU citizens, and due to the supranational nature of the EU jurisdiction, this ‘direct effect’ *does not depend on whether a MS takes a monist or a dualist approach to international law*. Even the United Kingdom, which has an outspoken dualist approach, had to accept that EU Regulations have direct effect within their national jurisdiction and might overrule Acts of Parliament. That is, as long as they were part of the EU.

The history of the EU goes back to the Second World War. In its aftermath, attempts were made to ensure economic interdependency between European states, thus hoping to contribute to the prevention of a new war. This led six states (the German Federal Republic (West Germany), France, Italy, the Netherlands, Belgium, and Luxembourg) to the institution of the European Coal and Steel Community (ECSC) in 1952, followed by the European Economic Community (EEC) in 1958, which aimed to institute a common internal market, enabled by ‘four freedoms’: the free movement of goods, persons, services, and capital. In 1992, the ECSC and the EEC were integrated into the EU, then comprising of twelve MSs.

For a long time, the main purpose of the EU was to *harmonize* the legislation and the policies of its MSs in order to prevent obstruction and disruption of the internal economic market. If the sale of washing machines is subject to different legal requirements in different MSs, it becomes more difficult for manufacturers and retailers to produce and sell such machines across national borders. The same goes for, for example, data protection legislation; if the constraints for the processing of personal data differ per MS, cross border data processing becomes a problem that will, for example, reduce cross-border eCommerce.

By now, the EU, comprising of twenty-seven states, has a broader objective than merely the creation and protection of an effective and efficient economic market, as it more explicitly targets instituting an area of freedom, security, and justice without internal frontiers. This is most visible in the enactment of the CFREU that came into force in 2009.³ The Charter addresses not only the institutions and bodies of the Union but also the MSs whenever they implement Union law, while providing fundamental rights to EU citizens.

4.3.1 Transfer of sovereignty

As one can imagine, a transfer of sovereignty implies a substantive and substantial interference with national sovereignty. The idea that MSs have transferred part of their sovereignty to a new entity with its own jurisdiction was consolidated in the case law of the highest court of the (then) EEC. In the seminal ‘Van Gend en Loos’ case of 1963,⁴ the highest Court of the EU, the Court of Justice of the European Union (CJEU) considered that:

The objective of the EEC Treaty, which is to establish a Common Market, the functioning of which is of direct concern to interested parties in the Community, implies that this Treaty is more than an agreement, which merely creates mutual obligations between the contracting states. This view is confirmed by the preamble to the Treaty, which refers not only to governments but also to peoples. It is also confirmed more specifically by the establishment of institutions endowed with sovereign rights, the exercise of which affects Member States and their citizens . . .

³ Charter of Fundamental Rights of the European Union 2012/C 326/02.

⁴ CJEU, 5 February 1963, Case 26-62.

This firmly establishes the transfer of specified sovereign rights, followed by a transformative ‘speech act’ that in fact declared and instituted the EEC as a supranational legal order:

The conclusion to be drawn from this is that the community constitutes a new legal order of international law for the benefit of which the states have limited their sovereign rights, albeit within limited fields, and the subjects of which comprise not only Member States but also their nationals.

This has consequences for the freedom of MSs with regard to accepting ‘direct effect’ within their national legal order, as explained in the seminal *Costa/ENEL* case of 1964,⁵ where the court states:

By contrast with ordinary international treaties, the EEC Treaty has created its own legal system which, on the entry into force of the treaty, became an integral part of the legal systems of the Member States, and which their courts are bound to apply (...). The executive force of community law cannot vary from one state to another in deference to subsequent domestic laws, without jeopardizing the attainment of the objectives of the treaty.

The consequences—indeed the legal effect—of this judgment can hardly be overestimated. In accepting this judgment, the MSs have accepted that the EEC, which is now the EU, constitutes *a legal order in its own right, with jurisdiction over aspects of the national legal orders of the MSs*.

It remains important to note that the Constitution of a MS must allow for the transfer of sovereign power. In the Netherlands Constitution, the competence for such transfer can be found in Article 92, referring back to the conditions stipulated in Article 91(3), which has been quoted above, states:

Legislative, executive, and judicial powers may be conferred on international institutions by or pursuant to a treaty, subject, where necessary, to the provisions of Article 91(3).

An arduous issue nevertheless remains: who determines the boundaries of EU legislative competence when a national constitutional court of a MS disagrees with the position taken by the CJEU? This has been coined as the

⁵ CJEU, 15 July 1964 Case 6-64.

issue of *Kompetenz-Kompetenz*, as it concerns the competence to decide on competence. On several occasions this issue has arisen, notably when the German Constitutional Court (GCC) was asked to decide the legislative competence of the EU regarding issues that may infringe the German Constitution. So far, even though the GCC claims the competence to decide on these issues,⁶ it has not invalidated any judgment of the CJEU.⁷ Clearly, the CJEU is of the opinion that it is the only authority on the competences of the EU. Thus, by avoiding a disagreement, the GCC has saved the day, since competition over competence between the two highest courts could initiate the disintegration of the EU.

4.3.2 Sources of EU law

In the context of the EU, lawyers speak of the so-called ‘*acquis*’ (French for what has been achieved, established). This is the body of common rights and obligations that is binding on all the MSs of the EU.

The ‘acquis’ is constantly evolving and comprises:

- the content, principles, and political objectives of the Treaties;
- legislation adopted pursuant to the Treaties;
- the case law of the Court of Justice;
- declarations and resolutions adopted by the Union;
- instruments under the Common Foreign and Security Policy;
- instruments under Justice and Home Affairs;
- international agreements concluded by the Community and those entered into by the MSs among themselves within the sphere of the Union’s activities.

We end this chapter with the presentation of two types of legislative instruments that are core to EU law, and feature prominently in the second part of this book (as they regulate, e.g. data protection law, cybercrime, and copyright). Article 288 of the Treaty of the Functioning of the European Union (TFEU) specifies:⁸

⁶ BVerfG, 30 June 2009, 2 BvE 2/08.

⁷ BVerfG, 7 September 2011, 2 BvR 987/10.

⁸ European Union, Consolidated version of the Treaty on the Functioning of the European Union, 13 December 2007, 2008/C 115/01, available at: <http://www.refworld.org/docid/4b17a07e2.html>.

To exercise the Union's competences, the institutions shall adopt regulations, directives, decisions, recommendations and opinions.

A regulation shall have general application. It shall be binding in its entirety and directly applicable in all Member States.

A directive shall be binding, as to the result to be achieved, upon each Member State to which it is addressed, but shall leave to the national authorities the choice of form and methods.

A decision shall be binding in its entirety. A decision which specifies those to whom it is addressed shall be binding only on them.

Recommendations and opinions shall have no binding force.

We focus on regulations and directives as legislative instruments.

Regulations have 'direct effect' in all the MSs, whether or not they take a dualist approach to international law. Regulations are: (1) of general application; (2) binding in their entirety; and (3) directly applicable. Note that the 'direct effect' is not a consequence of states following the monist approach, but a consequence of the supranational character of EU law.

Next to regulations, the EU has another type of legislative instrument, namely *directives*. Though they are binding law, they lack 'direct effect'. Instead they impose an obligation on the MSs to transpose the content of the directive into their own legal system, that is, they may have to amend existing legislation or enact new statutes. Directives basically dictate that certain results must be achieved, while leaving it to the MSs' discretion how to achieve this, depending on their own legal system and the legal culture it embeds.

The difference between regulations and directives marks the challenges of the EU, which is neither a superstate nor a form of collaboration based on international law. On the one hand, some legislation is formulated in one and the same way and applies unilaterally in all MSs (regulations). On the other hand, some legislation has to be adapted by the MSs, taking into account how it could best fit with and within their legal order (directives). The latter leaves more room for different uptake in the different MSs, which may be confusing for transnational players on the internal market, but will be better adapted to local circumstances. Nevertheless, even regulations may be interpreted differently across different national jurisdictions, thus jeopardizing the goals of harmonization that are core to the EU.

4.3.3 Case law of the CJEU

To prevent contradictory interpretations of EU law, courts in the MSs may consult the CJEU in a so-called ‘preliminary proceeding’, inviting the Court to provide an authoritative interpretation of EU law for the case at hand. Such preliminary rulings bind all the MSs and thus further the harmonization of law in the EU, including the harmonization of legal protection against violations of fundamental rights. In section 2.1.2.1, we have already encountered a landmark case of the CJEU on the validity of the Data Retention Directive, in the light of the CFREU. We can now understand the relevance of the fact that this concerned a directive, since directives must be implemented in national law. The Court’s judgment that declared the directive *invalid*, did not necessarily affect its national implementation. All MSs had to check whether their national law—based on the directive—complied with the relevant legal conditions identified by the Court for valid data retention duties for the telco operators. We reiterate these legal conditions as recounted in section 2.1.2. To qualify as *lawful restrictions of the rights to privacy and data protection*, measures enacted as an implementation of the Data Retention Directive must, even if they have a legitimate aim and are appropriate to achieve this aim, nevertheless be proportional.

According to the CJEU, this entails that:

- the measures are sufficiently circumscribed, limited to what is strictly necessary;
- the scope of the retention measures must be differentiated;
- relevant limitations and/or exceptions must be foreseen; as well as
- objective criteria to ensure that data is only used for the most serious offences;
- the retention period should differentiate between categories of data;
- storage outside the EU should be prohibited.

To assess whether the transposition of the directive complies with the Court’s interpretation, each MS had to check their legislation and policies against these criteria. In some MSs, the legislature found that they were compliant, whereas in other MSs, courts found the relevant transposition to be in violation of Articles 7 and 8 CFREU. In point of fact, two cases were referred to the CJEU, asking whether or not national transposition was in violation, notably *Tele2 Sverige AB v. Post- och telestyrelsen* and *Secretary of State for the Home*

Department v. Watson.⁹ In both cases the CJEU found that the national legislation was indeed in violation. The reasoning concerns the fact that such national legislation must comply with Article 15 of the ePrivacy Directive, which protects the confidentiality of electronic communication. Article 15 of the ePrivacy Directive allows MSs to restrict the applicability of some articles, based on national legislation, if such national legislation is restricted to the goals stipulated in Article 15, contains proper safeguards, and is necessary in a democratic society (the proportionality requirement).

Such proportionality, according to the CJEU, is absent in the case of national legislation:¹⁰

which, for the purpose of fighting crime, provides for general and indiscriminate retention of all traffic and location data of all subscribers and registered users relating to all means of electronic communication.

national legislation governing the protection and security of traffic and location data and, in particular, access of the competent national authorities to the retained data, where the objective pursued by that access, in the context of fighting crime, is not restricted solely to fighting serious crime, where access is not subject to prior review by a court or an independent administrative authority, and where there is no requirement that the data concerned should be retained within the European Union.

We should note that, as in the case of international treaties, the legislative instruments of the EU often contain a number of recitals, which are not legally binding in the way that articles are, but nevertheless pivotal for the interpretation of these articles. For instance, in the judgment of *Tele2 Sverige AB v Post- och telestyrelsen* and *Secretary of State for the Home Department v. Watson*, the Court states in paragraph 87:

The scope of Article 5, Article 6 and Article 9(1) of Directive 2002/58, which seek to ensure the confidentiality of communications and related data, and to minimise the risks of misuse, must moreover be assessed in the light of recital 30 of that directive, which states: ‘Systems for the provision of electronic communications networks and services should be designed to limit the amount of personal data necessary to a strict minimum’.

⁹ CJEU, 21 December 2016, Cases C-203/15 and C-698/15 (*Tele2 Sverige AB v Post- och telestyrelsen* and *Secretary of State for the Home Department v. Watson*).

¹⁰ Dictum (decision) CJEU, 21 December 2016, Cases C-203/15 and C-698/15, under 1 and 2.

And, in paragraph 95 the Court states that:

(...) As regards recital 11 of that directive, it states that a measure of that kind must be ‘strictly’ proportionate to the intended purpose. In relation to, in particular, the retention of data, the requirement laid down in the second sentence of Article 15(1) of that directive is that data should be retained ‘for a limited period’ and be ‘justified’ by reference to one of the objectives stated in the first sentence of Article 15(1) of that directive.

This clearly demonstrates how recitals may not be binding, but are indeed an important source of law.

4.4 International Rule of Law

International law depends on national law. First, because national law determines to what extent states are bound by international law. Second, because enforcement of international law depends on national bodies (legislature, courts, administration). This implies that international law, to a large extent, depends on states willing to bind themselves. There are some exceptions, for example, with regard to *ius cogens*, which applies whether or not states recognize its force. But, generally speaking, one may be tempted to assume that states act as legal subjects in the realm of international law, free to negotiate treaties and free to subject themselves to whatever they deem to be in their own interest.

However, national law also depends on international law. First, because the system of sovereign states is based on mutual recognition of each other’s internal and external sovereignty. As discussed in sections 1.4 and 4.1.2, sovereignty is an artificial construct, a historical artefact. Without external sovereignty, which depends on the international legal order, we cannot ‘have’ internal sovereignty. In the words of Jeremy Waldron:

In its municipal [national, mh] aspect, the state is a particular tissue of legal organization: it is the upshot of organizing certain rules of public life in a particular way. Its sovereignty is something made, not assumed, and it is made for the benefit of those whose interests it protects. In its international aspect, the sovereignty and sovereign freedom of the individual state is equally an artifact of international law. What its sovereignty is and what it amounts to is not given as a matter of the intrinsic value of its individuality, but determined by the rules of the international order.

This leads Waldron to quite another understanding than that of states as legal subjects that are free to act in their own interest. Instead he considers them as both sources of international law and *officials of international law*.

The latter implies that from the perspective of the rule of law, states are not free to act in their own interests but bound by a legality principle at the level of international law. In the context of national law, the rule of law means that citizens are not there to serve the state, but the state is there to serve its citizens. In the context of international law, the rule of law means that states serve as the trustees of their citizens, bound to the rule of international law not for the sake of their own sovereignty, but for the sake of the people whose well-being they are entrusted with. Because this fiduciary position of states depends on the international legal order, to some extent they are also officials of the international legal order. Ultimately, this may entail a responsibility of states for subjects of other states.

References

Introduction to international and supranational law

- Glahn, Gerhard von, and James Larry Taulbee. 2017. *Law Among Nations: An Introduction to Public International Law*. 11th ed. New York: Routledge.
- Hage, Jaap, Antonia Waltermann, and Bram Akkermans, eds. 2017. *Introduction to Law*. 2nd ed. New York: Springer, chapters 10 and 12.
- Schütze, Robert. 2012. *An Introduction to European Law*. Fourth Impression ed. Cambridge and New York: Cambridge University Press.
- Simpson, G., ed. 2001. *The Nature of International Law*. London: Routledge.

On the sources of international law

- Besson, Samantha, and Jean d'Aspremont. 2017. 'The Sources of International Law'. *The Oxford Handbook of the Sources of International Law*, October. <https://doi.org/10.1093/law/9780198745365.003.0001>.

On the rule of law in international law

- Waldron, Jeremy. 2006. 'The Rule of International Law'. *Harvard Journal of Law & Public Policy* 30 (1): 15–30 (quote at 21).

On *Kompetenz-Kompetenz*

Beck, Gunnar. 2011. "The Lisbon Judgment of the German Constitutional Court, the Primacy of EU Law and the Problem of *Kompetenz-Kompetenz*: A Conflict between Right and Right in Which There is No *Praetor*." *European Law Journal* 17 (4): 470–94. <https://doi.org/10.1111/j.1468-0386.2011.00559.x>.

PART II

DOMAINS OF CYBERLAW

Having provided a first introduction to ‘what law does’ and to ‘how it operates’ in Part I, we can now proceed to more specifically relevant legal domains in Part II. As our shared environment is increasingly ‘run’ by code- and data-driven systems, those who develop, sell, integrate, tweak, or employ them, as well as those who are subject to their automated decisions, need to confront human rights law, notably privacy and data protection; cybercrime law; copyright law; and private law liability for harm caused. This entails an inquiry into the relevant sources of law, notably legislation and case law, demonstrating more concretely how law and the rule of law operate in the era of a surging dependence on computational ICIs.

5

Privacy and Data Protection

Working with computing systems, whether developing, integrating, or testing them, will often involve working with data. Sometimes this data will be personal data, and sometimes these systems will have a major impact on the private life of those targeted by these systems (think of data brokers, credit rating agencies), or those interacting with these systems (in the case of social networks, search engines). In this chapter, we will investigate the legal domain of privacy and data protection, which entails a series of *legal requirements* for the development and design, for the default settings, and for the employment of computer architectures. This chapter can in no way provide a comprehensive overview of privacy and data protection, which would require two separate books at the least. However, the purpose of this book is not to turn computer scientists into lawyers. The purpose is to provide some real taste and true bite of the law on legal topics that are highly relevant for computer science. Therefore, please check the references for further reading and for real world scenarios check with a practising lawyer.

The right to privacy is a *subjective right*, attributed by *objective law*. This may be national (constitutional) law, international human rights law, or supranational law (EU fundamental rights law). In this chapter, we will first confront the landscape of human rights law at the global, national, and EU level, followed by a discussion of the concept of privacy. We will then inquire into the right of privacy, as guaranteed under the European Convention on Human Rights (ECHR) and the Charter of Fundamental Rights of the European Union (CFREU), and finally, we will target the new fundamental right to data protection, as guaranteed by the CFREU and protected by the General Data Protection Regulation (GDPR).

5.1 Human Rights Law

When tracing the history of human rights, we first encounter the English *Bill of Rights* of 1689, followed by the revolutionary French *Déclaration des Droits de l'Homme et du Citoyen* of 1789 and the US *Bill of Rights* of 1791. Though the famous *Magna Charta* of 1215 may seem an early example of a

human rights charter, it did not attribute what we now call human rights. Instead, it ensured that the feudal lords were able to restrict the powers of the King, while protecting jurisdiction over their own subjects against royal interference. The era of the *Magna Charta* saw the struggle between a feudal society and an emergent royal power; this was not yet the era of a powerful modern state that managed to subject each and every person on its territory to its jurisdiction. The rights provided by the *Magna Charta* were mainly reserved for powerful lords, who wished to preserve the powers they had over their own land and their own serfs against the claims of the king.

5.1.1 Human rights as defence rights against the modern state

The rise of the modern state must be situated in the beginning of what historians call the era of 'Modernity', around the fifteenth and sixteenth century. It was the rise of the modern, bureaucratic state that warranted new types of protection against the monopolistic powers of the King and his clerks (feeding on the impressive affordances of proliferating printed text, see section 1.4). The rise of the idea of human rights coincides with the rise of sovereignty (see section 1.4 and 4.1.2).

The human rights declarations of the seventeenth and eighteenth centuries provided those subject to the power of a sovereign state with an entitlement to civil and political rights, emulating their status to that of individual right bearers and constituents of the polity.

Being subject to a sovereign became being a subject in law. It is hard to imagine how novel the attribution of such individual, subjective rights was, even if initially their enforcement was neither practical nor effective.

Some attribute the power of this attribution to the 'endowment bias'; if people come to believe they 'have' these rights, they will invest in 'keeping' them. If the struggle this entails succeeds, these rights will eventually be instituted as effective subjective rights. In due course, respect for human dignity and a new emphasis on the centrality of the individual reconfigured the idea of law and politics, laying the groundwork for the more 'practical and effective' human rights protection of the second half of the twentieth century.

However, in the context of international law, human rights have been citizens' rights rather than human rights, depending on constitutional protection and citizenship, thus offering little protection for subjects of rogue states. After the atrocities of the Second World War, states decided to elevate the protection of human rights to the level of international law, starting with the *Universal Declaration of Human Rights* of 1948. Though this declaration had no binding force, it was soon followed by various treaties at the global and regional level, aiming to finally institute human rights as enforceable subjective rights against the state.

5.1.2 From liberty rights to social, economic, and further rights

Human rights law was originally focused on the protection of individual citizens against powerful states. We call these rights *first generation human rights*, and they are best described as the subjective right that the state refrains from interference with the legal good that is protected by such rights. This is why they are often called *liberty rights*.

These legal goods are: privacy, non-discrimination, bodily integrity, freedom of movement, the presumption of innocence, a fair trial, freedom of expression, freedom of association, freedom of religion, and voting rights. Note that these legal goods are considered worthy of protection as public goods, because a society that does not protect them cannot support a viable democracy that depends on independence of thought and unhindered development of both individual and group identities. For that reason, they are also called *civil and political rights*. The focus is on public goods that protect individual persons as autonomous agents in a democratic polity and on *negative obligations of the state* towards its citizens.

A *second generation of human rights* developed when it became clear that (1) non-interference is not always enough to protect such public goods, while (2) a number of other public goods were absent in the initial inventories of human rights. The public goods protected by second generation human rights concern public, for instance, employment, food and housing, social security, healthcare, and access to basic utilities such as electricity, postal services, and public transport.

These rights are often called *social and economic rights*. To actually provide these protected goods, a state cannot restrict itself to respecting liberty rights. The second-generation human rights impose *positive obligations* on states to create and sustain the goods it must protect. This implies that the second generation of human rights addresses states with ‘instruction norms’, rather than providing citizens with directly enforceable subjective rights. To exercise a right to employment, an economic system must be in place that enables such a right, meaning that second generation human rights require states to build institutions capable of supporting economic welfare and a fair distribution of access to social and economic goods.

Taking note that second generation human rights are instruction norms to states, rather than directly enforceable individual rights, the latter part of the twentieth century witnessed advocacy for a *third generation of human rights*.

Here, we encounter rights to construct and develop group identities and rights to a sustainable environment. These rights have even less of a straightforward relationship with individual entitlement, focusing on the *rights of groups* (e.g. the right to self-determination for indigenous peoples, which we already encountered in section 4.2.1, as a fundamental principle of international law) and *obligations towards the natural environment* on which human society depends (responsible innovation, sustainable development).

5.2 The Concept of Privacy

Before investigating the *right to privacy* as part of the first generation of human rights law, we will first inquire into *the nature of privacy itself*. The reason is that computer science has a specific relationship with privacy, notably in the context of digital security and cryptography. In that context, privacy is often seen as a subset of security, focused on hiding or removing the link between data and whoever the data refer to, or on encrypting the data to safeguard confidential data against eavesdropping. This has, as a consequence, meant that privacy protection is restricted to (1) anonymization or pseudonymization of personal data, by way of deleting or separating identifiers and to (2) hiding the content by means of encryption or other security measures. The focus on hiding has

generated research fields such as differential privacy and reidentification metrics, based on e.g. cryptography and key-management, k-anonymity, linkability metrics, and so on.

Though such research is of crucial importance to protect privacy, one must not mistake issues of identifiability and confidentiality for issues of privacy as the latter concerns far more than mere technical identifiability or readability.

Consider the following data points:

- your name;
- your bank account;
- the taxes your mother pays;
- what kind of socks you wear;
- the logs of your surfing behaviour on the net;
- your pattern of your energy usage behaviour;
- the decision to have an abortion;
- the decision, or inclination, to be a vegetarian.

Should we qualify this data as part of the privacy of the person the data refers to?

To answer this question, we need to check what falls within the *value of*, the *interest in*, or the *right to* privacy:

- When (under what conditions)?
- With regard to whom (is data on my mother part of my privacy)?
- Where (are specific locations more privacy-sensitive than others)?
- For what reason (what could make my socks relevant to my privacy)?

5.2.1 Taxonomies and family resemblance

Many authors have made attempts to define privacy by summing up the common denominators of what is generally seen as falling within the scope of privacy. This turns out to be a questionable undertaking, because the concept is as elusive as it is pertinent. Another way of tackling the issue of understanding privacy is to define it in terms of family resemblance.

The American privacy scholar and lawyer Daniel Solove made an insightful attempt to approximate the concept of privacy in terms of six categories that are partly overlapping, while thus covering much of what we intend when referring to privacy:

1. the right to be left alone;
2. limited access to self;
3. secrecy—concealment;
4. control over personal information;
5. personhood—protection of identity, dignity; and
6. intimacy.

Solove notes that some of these categories focus on goals, others on means, while they are in various way interdependent. Taken separately, none of these definitions would exhaust the concept of privacy, being either too broad or too narrow. He warns that this is therefore *not a taxonomy*, which would assume mutually independent features of the same thing. On the contrary, the idea of a *family resemblance* means that privacy cannot be defined in terms of necessary and sufficient conditions, because there is no common core to the different conceptions of privacy. Instead, Wittgenstein's notion of family resemblances enables us to take a pragmatic approach, recognizing the contextual, historical, dynamic nature of privacy, such as relating to family life, the body, or the home. This approach is bottom-up rather than abstract and acknowledges that, in the end, privacy is best seen as a set of practices rather than a formula. The concept of family resemblance was introduced as a way to understand the meaning of words by Wittgenstein in his *Philosophical Investigations*. The concept is very interesting for computer science as it explains why translating concepts into ontologies or a semantic web may entail a loss of meaning. I will therefore quote *The Stanford Encyclopedia of Philosophy* to elucidate this understanding of meaning:

There is no reason to look, as we have done traditionally—and dogmatically—for one, essential core in which the meaning of a word is located and which is, therefore, common to all uses of that word. We should, instead, travel with the word's uses through 'a complicated network of similarities overlapping and criss-crossing' (PI 66).¹ Family resemblance also serves to exhibit the lack of boundaries and the distance from exactness that characterize different uses of the same concept.

¹ This refers to para. 66 of Wittgenstein's *Philosophical Investigations*. See the correct reference to the *Stanford Encyclopedia* entry under references.

Such boundaries and exactness are the definitive traits of form—be it Platonic form, Aristotelian form, or the general form of a proposition adumbrated in the *Tractatus*.² It is from such forms that applications of concepts can be deduced, but this is precisely what Wittgenstein now eschews in favor of appeal to similarity of a kind with family resemblance.

To emphasize the elusive nature of privacy, we briefly follow Solove's discussion of the categories enumerated above.

A right to non-interference seems a pivotal shorthand for the right to privacy, as it clearly depicts the *negative obligations* of governments and others (vertical and horizontal effects of human rights law). Here, we think of privacy as the 'right to be left alone', where privacy is a liberty or freedom, in the sense of *freedom from external constraints*.

This understanding of privacy is related to *intimacy*, to the idea of drawing boundaries around a small circle of people with whom one dares to expose oneself, sharing information that might otherwise be used to shame a person, or to diminish or ridicule their agency. Intimacy relates to trust, not in the sense of confidence and security, but in the sense of trusting others enough to take the risk of being betrayed. One could ask what information is intimate, but this assumes that 'intimacy' is a property of information, whereas all depends on the situation, the context, and the roles played by intimate others. In some situations, financial information, or information shared with a health insurance company, may be intimate information, because it reveals to others what makes a person vulnerable to shame, ridicule, or even to life-threatening manipulation.

If we then take together privacy as *limited access*, and *secrecy, anonymity and solitude*, we can address the legal notion of third-party disclosure.

In the United States, the Supreme Court decided, in 1967,³ that once a person exposes their personal data to a third party such as banks or other service providers, they have no reasonable expectation of privacy regarding access

² The *Tractatus* is Wittgenstein's seminal work, preceding his *Philosophical Investigations*. In the latter, he rejects propositional logic and definitions in terms of sufficient and necessary reasons, though he endorsed them in the former. From the perspective of the latter, the view point taken in the former is just one 'language game' amongst many others, noting that the former should not claim a monopoly on understanding meaning.

³ *Katz v. United States*, 389 U.S. 347, 360 (1967), confirmed in, e.g. *California v. Greenwood*, 486 U.S. 35, 41 (1988).

by the government. This so-called ‘third-party doctrine’ reflects an approach to privacy that is radically different from the European approach, which does not presume that disclosing private information to one entity necessarily implies that other entities are now free to obtain and use such information.

Note that the United States have since enacted legislation requiring a warrant for access to specific data, thus providing specified protection for, for example, financial data and telephone data. We have already encountered the case of *US v. Jones* (followed by *Riley* and *Carpenter*, see section 2.1.2, n. 2), where the Supreme Court decided that police warrants were necessary in the case of GPS trackers, information on a cell phone, and cell-site records of a wireless carrier. These judgments may lead to the end of the third-party doctrine, depending on subsequent case law.

The next category, *control over information about oneself*, is often portrayed as the core meaning of what Americans call *informational privacy*. This understanding clearly links to the notion of identifiability, as it relates to information about an identifiable person, thus also connecting this particular conception of privacy with the idea of privacy as a subset of digital security.

Defining privacy in terms of control comes close to thinking of *personally identifiable information* (PII) as if it were the property of the person it concerns. PII is, just like informational privacy, a term used in the United States, whereas in the EU we generally speak of *data protection* and *personal data*. Thinking of PII in terms of property creates a number of problems, as neither data nor information are rivalrous or exclusionary. One person ‘having’ certain information does not necessarily imply that others do not ‘have’ that same information, whereas one person possessing a book implies that others do not possess it. It is therefore important to distinguish between control over ‘access to’ and ‘usage of’ information on the one hand, and property rights in information on the other. The latter applies in the case of intellectual property rights (e.g. copyright or patent), but not in the case of personal data. Below, we will discuss to what extent EU data protection law provides control to data subjects (those to whom personal data refers), but we can already point out here that full control over one’s personal data ignores the relational nature of personal data. To illustrate the latter point, we can think of Robinson Crusoe and ask the question whether he had a name before Friday came to his island. We have a name to be singled out by others, to be addressed by others, and to appear as a singular individual person before others. This implies that, though

we need some control over the sharing of our name, such control cannot be unlimited. Without fellows to address us, we effectively 'have' no name.

Finally, privacy is connected with *personhood*, with individuality, with dignity, and with autonomy. One could ask to what extent our personhood is private, noting that becoming a person depends on anticipating how others will frame us. Whereas the right to privacy is often seen as a liberty, as a right to be left alone, as a *freedom from* outside interference, privacy is also connected with a right to develop one's own identity, to be treated as worthy of respect, and the *freedom to* make one's own choices concerning, for example, lifestyle, employment, education, and political opinion. Here, privacy sits on the cusp of *freedom from* unreasonable constraint and the *freedom to* construct one's identity.

Indeed, this is how Agre and Rotenberg defined privacy, highlighting the interrelationship between negative and positive freedom. This also suggests that liberty and autonomy overlap and support each other. For instance, what has been called 'decisional privacy' (e.g. the right of a woman to decide about an abortion) clearly marks the nexus of *positive freedom* (to decide an abortion) with *negative freedom* (to be free from unreasonable constraints on such a decision). The crux of Agre and Rotenberg's definition resides in the requirement that people are free from *unreasonable* constraints, not just any constraints. In case law, legislation, and doctrine the concept of 'reasonable' or 'unreasonable' is of prime importance. Instead of framing this as a source of uncertainty, because of its prima facie vagueness, this concept can be seen as an aid in aligning different conceptions of legal goods that warrant protection. Demanding that a duty of care is exercised in a reasonable way acknowledges that 'a duty of care' cannot be defined in the abstract, but is better understood in terms of family resemblances. The duty of care of a mother, an employer, a manufacturer, and a social network provider may not share any common element; they nevertheless align along the lines of reasonable expectations and proper checks and balances, considering the relevant context and the roles of the parties involved. Similarly, reasonable expectations of privacy depend on context, on roles played, on checks and balances, and meaningful choice. This is not because privacy is a vague concept but because the practice of privacy is complex, requiring acuity to what is at stake for whom.

Though the reader may by now be wary of the dynamic and shifting borders of the concept of privacy, it is crucial to sustain awareness that privacy is a moving target.

Defining privacy in terms of *necessary and sufficient conditions* would restrict protection to what happens to fall within their scope, easily rendering the concept both *over- and under-inclusive*.

In the end, defining privacy is a decision to be taken when confronted with its violation. As Solove saliently writes in reference to a famous American philosopher:

‘[K]nowledge is an affair of making sure,’ Dewey observed, ‘not of grasping antecedently given sureties.’

This is what the courts must achieve every time a case is brought before them: *making the difference that makes a difference*.

5.2.2 Privacy and technology

After tracing the conceptual challenges of delineating privacy, I will briefly trace the relationship between privacy and technology. Some of us may think that privacy is a property of people in general, just like animals often display what ethologists call ‘critical distance’ from each other.

Privacy, according to environmental psychologist Altman, is a matter of shaping and negotiating borders between self and others. It is not a property of a person, but of a relationship.

Rather than being a matter of seclusion, Altman frames privacy as a continuous process of sharing and excluding, based on societal practices that are in turn dependent on technological affordances of the environment. In that sense, privacy can be detected in most human societies, though under different names and with very different constraints.

The right to privacy, however, is a recent historical artefact. As a subjective right, the right to privacy first surfaced at the end of the nineteenth century, in response to the proliferation of technologies such as photography and mass media.

In a famous article in the *Harvard Law Review*, US legal scholars Samuel Warren and Louis Brandeis discussed the need to protect oneself against

publication of photographs without permission, to enable *social withdrawal*. In that article, they formulated the right to privacy as the right to be left alone, basically arguing for the existence of a *privacy tort* whenever this right was infringed upon without justification. Interestingly, privacy was thus introduced as a private law issue rather than a constitutional right. When Brandeis later served as justice in the Supreme Court, however, he argued that such a right to be left alone must be ‘read into’ the US Constitution, notably into the Bill of Rights, thus vouching for a right to privacy against the state. The rise of mass media and photography afforded massive dissemination of pictures taken, thus infringing the privacy of those concerned in a previously unprecedented manner. This, in turn, gave rise to defence mechanisms to safeguard one’s capability to withdraw from such exposure.

This first appearance of a right to privacy fostered privacy as *negative freedom*: the right that others refrain from interference.

After the Second World War, a new technological infrastructure surfaced to enable and improve public administration, in the form of computerized databases. This resulted in the collection and storage of myriad data relating to identifiable citizens, enabling government agencies to better target their constituency and to engage in what would now be termed ‘evidence-based policy’. This, in turn, raised the question to whom this data belongs. In 1967, Alan Westin wrote a seminal work on *Privacy and Freedom*, taking a clear stand on the question of who should—by default—be capable of controlling access to data concerning individual persons. Privacy, he wrote, is:

the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others.

This concept of informational privacy, as *control over information*, informs much of the debate about privacy and data protection in our current age. It is interesting to note that it emerged in counterpoint to the rise of databases in public administration, as well as private enterprise. The fact that data was collected, sorted, and recorded, enabling retrieval as well as aggregation, gave rise to new types of transparency, and new types of threats to personal identity. This was related to the fact that in this era the data collected and stored was mostly stable data, allowing the mapping of both individuals and populations

in consistent and foreseeable way, without the kind of dynamic and unstructured big data capture that characterizes the current era.

This, second appearance of the right to privacy fosters privacy as a *positive freedom*: the freedom to determine how personal information is shared and used.

After the rise of the internet and the world wide web, combined with the capture of big data and data-driven techniques to infer new information, the need for a more complex and contextual right to privacy seems obvious. Negative freedom will not do, as data abounds and is captured beyond one's control on a permanent basis. For the same reason, positive freedom seems unattainable, as consent loses its meaning amidst the volume, variety, and velocity of data capture, storage, and use. A more practical and effective way of understanding privacy should therefore combine negative and positive freedom, while highlighting the relationship with identity-construction, not merely identification.

The definition of Agre and Rotenberg, referred to above, may be the most apt for the era of proactive and pre-emptive computing infrastructures, depicting the right to privacy as:

the right to be free of unreasonable constraints on the building of one's identity.

For some readers, this may sound overly vague or complicated. To confront a complex, volatile, invasive, and pre-emptive environment we will, however, need an understanding of privacy that goes beyond the hiding of personal data.

5.3 The Right to Privacy

Privacy is a *value*, an *interest*, a *right*, or a *good*. It can be analysed from an ethical perspective (as a *value*, a *virtue*, or *duty*), from an economic perspective (as a *utility*, a *preference*, or an *interest*), and from the perspective of political theory (as a *public* and a *private good*). In this work, we will focus on the legal perspective, tracing positive law's applicability to issues of privacy. Below, we will discuss the right to privacy from the perspectives of constitutional, international, and supranational law, ending with a discussion of Article 8 ECHR.

5.3.1 The right to privacy: constitutional law

The right to privacy is a subjective right, attributed by objective law. The most obvious branch of objective law that attributes the subjective right of privacy is *constitutional law*, which often contains a section that aims to protect citizens against overly invasive powers of the state. Historically, human rights initially played out in the vertical relationship between state and citizens, not in the horizontal relations between private parties. The industrial revolution of the nineteenth century gave rise to powerful economic actors whose ability to infringe privacy, freedom of information, and non-discrimination increasingly matched the powers of the state.

This has led courts to recognize a so-called 'horizontal effect' of constitutional rights such as privacy. This entails that protection against such infringements is a duty of the state, meaning that citizens can sue the state for failing to impose prohibitions to infringe these rights upon powerful players in the private sector. This is called *indirect horizontal effect*, because it cannot be invoked directly against private parties.

Depending on national jurisdiction, courts may also attribute *direct horizontal effect*, when qualifying a violation of privacy by, for example, a company as a tortious act in the context of private law. In that case, violation of privacy can be invoked directly against, for example, a private company.

In many states outside the Council of Europe, the Constitution provides the main protection against infringements of the right to privacy. For instance, in the United States, even though neither the 1787 US Constitution nor the 1791 Amendments to the US Constitution (known as the *Bill of Rights*) explicitly refer to a right to privacy, the Supreme Court of the United States has nevertheless interpreted various articles of the *Bill of Rights* as safeguarding an individual right to privacy,⁴ notably based on the Fourth Amendment:

The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

⁴ First relevant case was *Griswold v. Connecticut*, 381 U.S. 479 (1965).

This Amendment protects against:

- ‘unreasonable searches and seizures’ by the police,
- which require ‘a warrant’,
- that may only be issued in the case of probable cause (concrete and objectifiable suspicion), and
- must contain a reasonably detailed description of what may be searched or seized.

We can read these protections in terms of legal conditions and legal effect, by stating that ‘searches and seizures’ by government officials are only lawful if:

- there is probable cause,
- a warrant has been issued,
- which contains limitations as to what is allowed.

As we have already seen in section 2.1.2 and 5.1.2, the question here is (1) whether this right protects against violation of property rights (trespass) or also against violation of reasonable expectations of privacy that do not depend on property and (2) whether search and seizure of, for example, a mobile phone falls within the scope of the Fourth Amendment, as a phone is neither part of a person, a house, paper, or effects.

In the United States, constitutional protection of the right to privacy (which is also ‘read into’ other parts of the *Bill of Rights*) thus depends on national law, rather than international law. This has consequences for its applicability in the case of those who have no legal status in the United States, as it may be unclear whether the *Bill of Rights* even applies to them. Another consequence is that the enforcement of rights against the state is dependent on that same state. In contrast, the ECHR offers a more layered architecture of legal protection, which is at least in part dependent on a European court that is not part of the state against which it aims to protect.

5.3.2 The right to privacy: international law

Protection of human rights requires a resilient system of checks and balances, that is, a series of institutional safeguards to ensure that the state does not claim unreasonable exceptions and faces a stringently independent judiciary

to keep the powers of the state ‘in check’. As noted above, the need to protect subjects of the state against the state, gave rise to *international human rights law*, which provides an extra layer of checks and balances. Privacy is explicitly protected by Article 17 of the United Nations (UN) International Covenant on Civil and Political Rights (ICCPR) of 1966, and by Article 8 ECHR of 1950, two examples of international law. Both articles are similar, we quote Article 8 ECHR to give the reader a first taste:

1. Everyone has the right to respect for his private and family life, his home and his correspondence.
2. There shall be no interference by a public authority with the exercise of this right except such as is in accordance with the law and is necessary in a democratic society in the interests of national security, public safety or the economic well being of the country, for the prevention of disorder or crime, for the protection of health or morals, or for the protection of the rights and freedoms of others.

The UN ICCPR has global application, with currently 178 signatories and 172 ratifications, but its enforcement mechanisms are relatively weak compared to the ECHR. In Article 34, the ECHR provides citizens of the forty-eight contracting parties with an individual right to complain to the European Court of Human Rights (ECtHR):

The Court may receive applications from any person, non-governmental organisation or group of individuals claiming to be the victim of a violation by one of the High Contracting Parties of the rights set forth in the Convention or the protocols thereto. The High Contracting Parties undertake not to hinder in any way the effective exercise of this right.

The ECHR, however, does not have global application, as it only applies within the jurisdiction of the Council of Europe.

5.3.3 The right to privacy: supranational law

Since 2009, when the CFREU came into force, the protection of human rights has gained even more traction, adding a second European Court with competence to test legislation, decisions, and actions against a catalogue of human rights. This protection, offered at the level of supranational law, is applicable whenever member states (MSs) ‘are implementing Union law’ (Article 51

CFREU). As human rights developed with the rise of the modern state, they further developed with the rise of supranational jurisdiction. The prevailing powers of the institutions of the EU demand countervailing powers in the form of supranational fundamental rights.

5.3.4 Article 8 ECHR

In this section, we will discuss one of the most crucial legal rights of this book. The right to privacy that is articulated in Article 8 ECHR is not only relevant for bodily integrity, decisional privacy, and the other aspects of privacy, but also directly affects issues of cybercrime and copyright. This is due to the fact that cybercrimes may violate privacy (hacking, data breaches), or that copyright holders may violate privacy when disseminating their works (photographs, texts), but also because the investigative measures that aim to detect cybercrime and violations of copyright often infringe upon the right to privacy as protected in Article 8.

Here, we develop a first analysis of the *legal conditions* stipulated by art. 8 ECHR, how they are explained by the ECtHR, and the *legal effects* they generate.

Article 8 consists of two paragraphs. The first paragraph concerns the question of whether privacy is infringed, the second paragraph clarifies under what conditions an infringement is justified.

1. Everyone has the right to respect for his private and family life, his home and his correspondence.

The legal effect generated by this paragraph is ‘an infringement of privacy’, and this infringement depends on the following *alternative* legal conditions:

- private life is not respected;
- family life is not respected;
- the protection of one’s home is not respected; and
- the confidentiality of one’s correspondence is not respected.

The ECtHR takes the view that these concepts require a broad rather than a narrow interpretation, bringing a wide variety of situations, events, relationships, and contexts under the protection of Article 8.

Private life can be at stake in the context of work, meaning that a search of an office space may be an infringement of privacy.⁵ Family life is at stake when a state prohibits members of a family from living together, for instance in the case of a refusal to provide a residence permit for a partner from another state, or of a parent wishing to further develop a relationship with their child despite not being married to the other parent. Protection of the home may become relevant when a person has taken residence in a house they neither own nor rent, meaning that the need to respect one's home is not dependent upon ownership or contract. The confidentiality of communication has been interpreted to include letters, telephone calls, and more recently all types of internet-enabled communication that is not public. Privacy, as protected by Article 8, clearly concerns physical, spatial, contextual, decisional, communicative, and informational privacy, and although Article 8 addresses the contracting states, its indirect horizontal effect has been recognized by the ECtHR, requiring states to ensure proper protection against violations by others than the state. Note that the individual complaint right of the ECHR can only be invoked against a state, not against a company. To invoke direct horizontal effect, a person needs to sue the tortfeasor in a national court.

An infringement of privacy is not the same as a violation of the right to privacy. Once the legal effect of an infringement has been established by the ECtHR, it will investigate whether the state has a valid justification.

2. There shall be no interference by a public authority with the exercise of this right except such as is in accordance with the law and is necessary in a democratic society in the interests of national security, public safety or the economic well being of the country, for the prevention of disorder or crime, for the protection of health or morals, or for the protection of the rights and freedoms of others.

The legal effect of a valid justification is that, despite the infringement, Article 8 is not violated. This effect depends on the following *cumulative* legal conditions:

- the infringement has one of more of the following legitimate aims: national security, public safety, or the economic well-being of the country, for the

⁵ ECtHR, 16 December 1992, Application no. 13710/88 (*Niemietz v. Germany*), regarding the search of a law firm; ECtHR, 25 June 1997, Application no. 20605/92 (*Halford v. UK*), regarding the interception of telephone calls at work.

prevention of disorder or crime, for the protection of health or morals, or for the protection of the rights and freedoms of others;

- the infringement is in accordance with the law; and
- the infringement is necessary in a democratic society.

The second paragraph of Article 8 thus requires a *triple test*, meaning that all three legal conditions must be met. These conditions are often summed up by stating that any infringing measures taken by the state must:

- have a legitimate aim;
- have a basis in law; and
- be proportional in relation to the aim served.

The articulation of *legitimate aims* in Article 8.2 is rather inclusive, which means that the ECtHR seldom finds reason to endorse the claim that the state lacked a legitimate aim.

Many of the cases where the ECtHR (the Court) finds that Article 8 has been violated concern the legal condition that the infringement must be ‘in accordance with the law’ to be justified. This basically refers to the legality principle of constitutional law (see section 3.3).

The Court has developed—over the course of the years—another *triple test* to decide whether an infringement has a proper basis in law:

- the legal competence to take infringing measures must be *accessible*, knowable for citizens to whom it will apply;
- the infringements must be *foreseeable*, which means sufficiently specified; and
- the quality of the law must include *sufficient safeguards* that limit the exercise of the competence in time and space, specifying the extent to which privacy may be infringed, and notably requiring independent oversight (e.g. warrants) in the case of more serious infringements.

Note that the Court will not merely check legislative or regulatory provisions, but test practical arrangements and actual safeguards to establish whether the infringing measures were taken ‘in accordance with the law’. Throughout its case law, the ECtHR demands that the rights attributed in the ECHR are both ‘practical and effective’, stating that:⁶

⁶ *Airey v. Ireland*, 9 October 1979, Series A, no. 32, para. 24.

[t]he Convention is intended to guarantee not rights that are theoretical or illusory but rights that are practical and effective (...).

If privacy is infringed with a legitimate aim, based on a legal competence that is accessible, foreseeable, while having sufficient safeguards, the final test is a *proportionality test*.

The proportionality test entails that the ECtHR investigates whether the measure was necessary in a democratic society, which requires—according to the Court—a *pressing social need* to resort to such measures.

- Under this criterion the Court will examine the gravity, invasiveness, and seriousness of the infringement in relation to the importance and seriousness of the aim served.
- This criterion basically requires that the measures taken can reasonably be expected to be effective, because a measure that is not effective cannot be necessary.
- The proportionality test includes a *subsidiarity test*; if another measure which is less infringing is feasible or sufficiently effective, the measure is not proportional.

5.3.5 Case law Article 8 ECHR regarding surveillance

When developing computing architectures, whether in the context of databases, streaming data, machine-to-machine communication, knowledge discovery in databases, machine learning, or cryptographic infrastructures, computer scientists lay the foundations for the ICIs that enable the processing, storage, interlinking, and inferencing of behavioural and other personal data. This may regard online clickstream behaviour, location, and mobility data, energy usage behaviours, biometric gait behaviour, and a plethora of communication data, including both content and metadata. Governments, tasked with the investigation and prosecution of criminal offences and the protection of national and public security, have many incentives to gain access to such data. Apart from the struggle against serious crime and threats to national security, governments need to collect taxes, attribute social benefits, take precautionary measures regarding public health, and safeguard the economic welfare of the country. All these tasks fall within the scope of the legitimate aims enumerated in Article 8.2 ECHR. This raises the question under what

conditions surveillance measures can be qualified as ‘in accordance with the law’ and if so, when they are considered ‘proportional’ to the targeted aim.

Surveillance measures by the police may regard post-crime investigatory measures (to identify an offender after a crime has been committed) or pre-crime investigations (to prevent potential offending, or to foresee likely offences). To understand how the Court deals with various types of electronic surveillance, we will discuss two cases of post-crime surveillance and two cases of pre-crime surveillance (including surveillance by the intelligence services, which falls outside the domain of criminal law).

This entails extensive quotation of the relevant case law, to *show* how the Court reasons, taking into account that the Court’s judgments bind the contracting parties and thus provide ‘practical and effective’ legal protection to those under the jurisdiction of the ECHR.

5.3.5.1 Post-crime surveillance

In 1984, in *Malone v. UK*,⁷ the ECtHR determined that the United Kingdom was in breach of Article 8 ECHR, where it allowed the interception of telephone conversations by the police upon a warrant issued by the Secretary of State. The Court determined that for such a measure to be ‘in accordance with the law’, it must not merely have a basis in domestic law (meaning a legal power), but must also be foreseeable and sufficiently limited as required by the rule of law:

68. Since the implementation in practice of measures of secret surveillance of communications is not open to scrutiny by the individuals concerned or the public at large, it would be contrary to the rule of law for the legal discretion granted to the executive to be expressed in terms of an unfettered power. Consequently, the law must indicate the scope of any such discretion conferred on the competent authorities and the manner of its exercise with sufficient clarity (...).

When applying this interpretation, the Court finds that:

79. The foregoing considerations disclose that, at the very least, in its present state the law in England and Wales governing interception of communications for police purposes is somewhat obscure and open to differing interpretations. The Court would be usurping the function of the national courts were it to attempt to make

⁷ ECtHR, 2 August 1984, Application no. 8691/79 (*Malone v. UK*).

an authoritative statement on such issues of domestic law (see, *mutatis mutandis*, the Deweer judgment of 27 February 1980, Series A no. 35, p. 28, in fine, and the Van Droogenbroeck judgment of 24 June 1982, Series A no. 50, p. 30, fourth subparagraph). The Court is, however, required under the Convention to determine whether, for the purposes of paragraph 2 of Article 8 (art. 8-2), the relevant law lays down with reasonable clarity the essential elements of the authorities' powers in this domain.

Detailed procedures concerning interception of communications on behalf of the police in England and Wales do exist (see paragraphs 42–49, 51–52 and 54–55 above). What is more, published statistics show the efficacy of those procedures in keeping the number of warrants granted relatively low, especially when compared with the rising number of indictable crimes committed and telephones installed (see paragraph 53 above). The public have been made aware of the applicable arrangements and principles through publication of the Birkett report and the White Paper and through statements by responsible Ministers in Parliament (see paragraphs 21, 37–38, 41, 43 and 54 above).

Nonetheless, on the evidence before the Court, it cannot be said with any reasonable certainty what elements of the powers to intercept are incorporated in legal rules and what elements remain within the discretion of the executive. In view of the attendant obscurity and uncertainty as to the state of the law in this essential respect, the Court cannot but reach a similar conclusion to that of the Commission. In the opinion of the Court, the law of England and Wales does not indicate with reasonable clarity the scope and manner of exercise of the relevant discretion conferred on the public authorities. To that extent, the minimum degree of legal protection to which citizens are entitled under the rule of law in a democratic society is lacking.

(iii) Conclusion

80. In sum, as far as interception of communications is concerned, the interferences with the applicant's right under Article 8 (art. 8) to respect for his private life and correspondence (see paragraph 64 above) were not 'in accordance with the law'.

In this case, Malone not only claimed that the interception of the content of his telephone conversations violated his right to privacy under the Convention, but also that the capture of what we would now call metadata violated said right. The Court states, with regard to this capture, known as 'metering':

83. The process known as 'metering' involves the use of a device (a meter check printer) which registers the numbers dialled on a particular telephone and the time

and duration of each call (see paragraph 56 above). In making such records, the Post Office—now British Telecommunications—makes use only of signals sent to itself as the provider of the telephone service and does not monitor or intercept telephone conversations at all. From this, the Government drew the conclusion that metering, in contrast to interception of communications, does not entail interference with any right guaranteed by Article 8 (art. 8).

87. Section 80 of the Post Office Act 1969 has never been applied so as to ‘require’ the Post Office, pursuant to a warrant of the Secretary of State, to make available to the police in connection with the investigation of crime information obtained from metering. On the other hand, no rule of domestic law makes it unlawful for the Post Office voluntarily to comply with a request from the police to make and supply records of metering (see paragraph 56 above). The practice described above, including the limitative conditions as to when the information may be provided, has been made public in answer to parliamentary questions (*ibid.*). However, on the evidence adduced before the Court, apart from the simple absence of prohibition, there would appear to be no legal rules concerning the scope and manner of exercise of the discretion enjoyed by the public authorities. Consequently, although lawful in terms of domestic law, the interference resulting from the existence of the practice in question was not ‘in accordance with the law’, within the meaning of paragraph 2 of Article 8 (art. 8-2) (see paragraphs 66 to 68 above).

Note that the ECtHR established that the practice of ‘metering’ is lawful under UK law, but in violation of Article 8.2 ECHR. Both the interception and the metering violate Article 8.2 because they are not ‘in accordance with the law’ as required by a treaty that binds the United Kingdom. This means that the United Kingdom has violated its legal obligations under the Convention and is now bound to ensure that these types of surveillance measures are based on a domestic law that both constitutes and sufficiently restricts its legal powers.

In 1990, in *Huvig & Kruslin v. France*,⁸ the ECtHR determined that Article 8 was breached. The case concerned the interception of telephone conversations, as in the *Malone* case. The Court extensively refers to its contentions in the *Malone* judgment as to the requirement of such interceptions being ‘in accordance with the law’. It then states:

35. Above all, the system does not for the time being afford adequate safeguards against various possible abuses. For example, the categories of people liable to

⁸ ECtHR, 24 April 1990, Application no. 11801/85 (*Huvig & Kruslin v. France*).

have their telephones tapped by judicial order and the nature of the offences which may give rise to such an order are nowhere defined. Nothing obliges a judge to set a limit on the duration of telephone tapping. Similarly unspecified are the procedure for drawing up the summary reports containing intercepted conversations; the precautions to be taken in order to communicate the recordings intact and in their entirety for possible inspection by the judge (who can hardly verify the number and length of the original tapes on the spot) and by the defence; and the circumstances in which recordings may or must be erased or the tapes be destroyed, in particular where an accused has been discharged by an investigating judge or acquitted by a court. The information provided by the Government on these various points shows at best the existence of a practice, but a practice lacking the necessary regulatory control in the absence of legislation or case-law.

36. In short, French law, written and unwritten, does not indicate with reasonable clarity the scope and manner of exercise of the relevant discretion conferred on the public authorities. This was truer still at the material time, so that Mr Kruslin did not enjoy the minimum degree of protection to which citizens are entitled under the rule of law in a democratic society (see the *Malone* judgment previously cited, Series A no. 82, p. 36, § 79). There has therefore been a breach of Article 8 (art. 8) of the Convention.

Note that in the *Huvig & Kruslin* judgment, the Court further details the nature of the restrictions that must be laid down by law, compared to the more general formulation in the *Malone* judgment.

5.3.5.2 Pre-crime surveillance (including surveillance by the intelligence services)

In 1978, in *Klass v. Germany*,⁹ the ECtHR decided a case regarding surveillance measures taken by the secret services in Germany. I will quote the most relevant considerations from the judgment, which should clarify *how the Court argues points of law* and thus *shapes the interpretation of legal conditions*:

All five applicants claim that Article 10 para. 2 of the Basic Law (Grundgesetz) and a statute enacted in pursuance of that provision, namely the Act of 13 August 1968 on Restrictions on the Secrecy of the Mail, Post and Telecommunications (... hereinafter referred to as 'the G 10'), are contrary to the Convention.

They do not dispute that the State has the right to have recourse to the surveillance measures contemplated by the legislation; they challenge this legislation in that it

⁹ ECHR, 6 September 1978, Application no. 5029/71 (*Klass v. Germany*).

permits those measures without obliging the authorities in every case to notify the persons concerned after the event, and in that it excludes any remedy before the courts against the ordering and execution of such measures.

Their application is directed against the legislation as modified and interpreted by the Federal Constitutional Court (Bundesverfassungsgericht).

The Court first discusses the admissibility of the complaint, raising the question whether the applicant is a victim of violation by one of the MSs.

33. (...) Article 25 (art. 25) [now Article 34, mh] does not institute for individuals a kind of *actio popularis* for the interpretation of the Convention; it does not permit individuals to complain against a law in abstracto simply because they feel that it contravenes the Convention. In principle, it does not suffice for an individual applicant to claim that the mere existence of a law violates his rights under the Convention; it is necessary that the law should have been applied to his detriment.

34. (...) The question arises in the present proceedings whether an individual is to be deprived of the opportunity of lodging an application with the Commission because, owing to the secrecy of the measures objected to, he cannot point to any concrete measure specifically affecting him. (...)

36. The Court points out that where a State institutes secret surveillance the existence of which remains unknown to the persons being controlled, with the effect that the surveillance remains unchallengeable, Article 8 (art. 8) could to a large extent be reduced to a nullity. It is possible in such a situation for an individual to be treated in a manner contrary to Article 8 (art. 8), or even to be deprived of the right granted by that Article (art. 8), without his being aware of it and therefore without being able to obtain a remedy either at the national level or before the Convention institutions. (...) The Court finds it unacceptable that the assurance of the enjoyment of a right guaranteed by the Convention could be thus removed by the simple fact that the person concerned is kept unaware of its violation. (...)

38. Having regard to the specific circumstances of the present case, the Court concludes that each of the applicants is entitled to '(claim) to be the victim of a violation' of the Convention, even though he is not able to allege in support of his application that he has been subject to a concrete measure of surveillance.

This entails that the Court makes an exception to the requirement that applicants must claim and demonstrate to be a victim of violation in concrete terms. Depending on the specific circumstances of the case at hand, the Court may decide to conduct an abstract test of relevant legislation, attributing the

status of ‘victims’ of what is now Article 34 ECHR, to *those who may have been a victim* of secret surveillance measures.

The Court then quotes relevant legislation, notably Article 10 of the Basic Law of Germany:

- (1) Secrecy of the mail, post and telecommunications shall be inviolable.
- (2) Restrictions may be ordered only pursuant to a statute. Where such restrictions are intended to protect the free democratic constitutional order or the existence or security of the Federation or of a Land, the statute may provide that the person concerned shall not be notified of the restriction and that legal remedy through the courts shall be replaced by a system of scrutiny by agencies and auxiliary agencies appointed by the people’s elected representatives.

The Court begins by investigating whether the legislation that is contested by the applicants, constitutes an *interference* with Article 8.1 ECHR:

41. The first matter to be decided is whether and, if so, in what respect the contested legislation, in permitting the above-mentioned measures of surveillance, constitutes an interference with the exercise of the right guaranteed to the applicants under Article 8 para. 1 (art. 8-1). (...)

Furthermore, in the mere existence of the legislation itself there is involved, for all those to whom the legislation could be applied, a menace of surveillance; this menace necessarily strikes at freedom of communication between users of the postal and telecommunication services and thereby constitutes an ‘interference by a public authority’ with the exercise of the applicants’ right to respect for private and family life and for correspondence.

As is often the case, the Court takes a broad view of the scope of the first paragraph and decides that the legislation constitutes an infringement. The next question is whether the infringement is justified:

42. The cardinal issue arising under Article 8 (art. 8) in the present case is whether the interference so found is justified by the terms of paragraph 2 of the Article (art. 8-2).

The Court first tests whether the infringement is ‘in accordance with the law’:

43. In order for the ‘interference’ established above not to infringe Article 8 (art. 8), it must, according to paragraph 2 (art. 8-2), first of all have been ‘in accordance with the law’.

This requirement is fulfilled in the present case since the ‘interference’ results from Acts passed by Parliament, including one Act which was modified by the Federal Constitutional Court, in the exercise of its jurisdiction, by its judgment of 15 December 1970 (see paragraph 11 above).

In addition, the Court observes that, as both the Government and the Commission pointed out, any individual measure of surveillance has to comply with the strict conditions and procedures laid down in the legislation itself.

This leads the Court to test whether the interference has a *legitimate aim*:

45. The G 10 defines precisely, and thereby limits, the purposes for which the restrictive measures may be imposed. It provides that, in order to protect against ‘imminent dangers’ threatening ‘the free democratic constitutional order’, ‘the existence or security of the Federation or of a Land’, ‘the security of the (allied) armed forces’ stationed on the territory of the Republic or the security of ‘the troops of one of the Three Powers stationed in the Land of Berlin’, the responsible authorities may authorise the restrictions referred to above (see paragraph 17).

46. The Court, sharing the view of the Government and the Commission, finds that the aim of the G 10 is indeed to safeguard national security and/or to prevent disorder or crime in pursuance of Article 8 para. 2 (art. 8-2). In these circumstances, the Court does not deem it necessary to decide whether the further purposes cited by the Government are also relevant.

This brings the Court to test the final criterion of the triple test, investigating whether the interference is necessary in a democratic society. Below you will find an extensive quotation of (part) of the reasoning of the Court regarding the question whether the interference enabled by the legislation is *proportional*, considering what is at stake.

47. The applicants do not object to the German legislation in that it provides for wide-ranging powers of surveillance; they accept such powers, and the resultant encroachment upon the right guaranteed by Article 8 para. 1 (art. 8-1), as being a necessary means of defence for the protection of the democratic State.

The applicants consider, however, that paragraph 2 of Article 8 (art. 8-2) lays down for such powers certain limits which have to be respected in a democratic society in order to ensure that the society does not slide imperceptibly towards totalitarianism. In their view, the contested legislation lacks adequate safeguards against possible abuse.

49. As concerns the fixing of the conditions under which the system of surveillance is to be operated, the Court points out that the domestic legislature enjoys a certain discretion. It is certainly not for the Court to substitute for the assessment of the national authorities any other assessment of what might be the best policy in this field (...)

Nevertheless, the Court stresses that this does not mean that the Contracting States enjoy an unlimited discretion to subject persons within their jurisdiction to secret surveillance. The Court, being aware of the danger such a law poses of undermining or even destroying democracy on the ground of defending it, affirms that the Contracting States may not, in the name of the struggle against espionage and terrorism, adopt whatever measures they deem appropriate.

51. According to the G 10, a series of limitative conditions have to be satisfied before a surveillance measure can be imposed. (...)

52. The G 10 also lays down strict conditions with regard to the implementation of the surveillance measures and to the processing of the information thereby obtained. (...)

53. Under the G 10, while recourse to the courts in respect of the ordering and implementation of measures of surveillance is excluded, subsequent control or review is provided instead, in accordance with Article 10 para. 2 of the Basic Law, by two bodies appointed by the people's elected representatives, namely, the Parliamentary Board and the G 10 Commission. (...)

54. The Government maintain that Article 8 para. 2 (art. 8-2) does not require judicial control of secret surveillance and that the system of review established under the G 10 does effectively protect the rights of the individual. The applicants, on the other hand, qualify this system as a 'form of political control', inadequate in comparison with the principle of judicial control which ought to prevail.

It therefore has to be determined whether the procedures for supervising the ordering and implementation of the restrictive measures are such as to keep the 'interference' resulting from the contested legislation to what is 'necessary in a democratic society'.

55. Review of surveillance may intervene at three stages: when the surveillance is first ordered, while it is being carried out, or after it has been terminated. As regards the first two stages, the very nature and logic of secret surveillance dictate that not only the surveillance itself but also the accompanying review should be effected without the individual's knowledge.

Consequently, since the individual will necessarily be prevented from seeking an effective remedy of his own accord or from taking a direct part in any review proceedings, it is essential that the procedures established should themselves provide adequate and equivalent guarantees safeguarding the individual's rights.

In addition, the values of a democratic society must be followed as faithfully as possible in the supervisory procedures if the bounds of necessity, within the meaning of Article 8 para. 2 (art. 8-2), are not to be exceeded.

One of the fundamental principles of a democratic society is the rule of law, which is expressly referred to in the Preamble to the Convention (see the Golder judgment of 21 February 1975, Series A no. 18, pp. 16–17, para. 34). The rule of law implies, *inter alia*, that an interference by the executive authorities with an individual's rights should be subject to an effective control which should normally be assured by the judiciary, at least in the last resort, judicial control offering the best guarantees of independence, impartiality and a proper procedure.

56. The Court considers that, in a field where abuse is potentially so easy in individual cases and could have such harmful consequences for democratic society as a whole, it is in principle desirable to entrust supervisory control to a judge.

Nevertheless, having regard to the nature of the supervisory and other safeguards provided for by the G 10, the Court concludes that the exclusion of judicial control does not exceed the limits of what may be deemed necessary in a democratic society.

58. In the opinion of the Court, it has to be ascertained whether it is even feasible in practice to require subsequent notification in all cases.

The activity or danger against which a particular series of surveillance measures is directed may continue for years, even decades, after the suspension of those measures.

Subsequent notification to each individual affected by a suspended measure might well jeopardise the long-term purpose that originally prompted the surveillance. Furthermore, as the Federal Constitutional Court rightly observed, such notification might serve to reveal the working methods and fields of operation of the intelligence services and even possibly to identify their agents.

In the Court's view, in so far as the 'interference' resulting from the contested legislation is in principle justified under Article 8 para. 2 (art. 8-2) (see paragraph 48 above), the fact of not informing the individual once surveillance has ceased cannot itself be incompatible with this provision since it is this very fact which ensures the efficacy of the 'interference'.

For these reasons the Court

1. holds unanimously that it has jurisdiction to rule on the question whether the applicants can claim to be victims within the meaning of Article 25 (art. 25) of the Convention;
2. holds unanimously that the applicants can claim to be victims within the meaning of the aforesaid Article (art. 25);
3. holds unanimously that there has been no breach of Article 8, Article 13 or Article 6 (art. 8, art. 13, art. 6) of the Convention.

This extensive quotation should contribute to a better understanding of the delicate and complex nature of the issues brought before the Court. This particular case (*Klass*) is a landmark case that functions as a building block for the reasoning in similar cases and requires the contracting states to incorporate necessary safeguards when developing and implementing legislation that enables surveillance by intelligence agencies.

In 2006, the ECtHR decided the case of *Weber & Saravia v. Germany*,¹⁰ once again testing legislation regarding so-called ‘strategic monitoring’ by intelligence services. In this case, the Court specifies in more detail what qualifies as ‘interferences’ that are ‘in accordance with the law’. Although, after having conducted the triple test, the Court decided that the contested legislation did not violate Article 8 ECHR, I will quote the legal conditions summed up by the Court to attain the legal effect of such interferences qualifying as being ‘in accordance with the law’.

95. In its case-law on secret measures of surveillance, the Court has developed the following minimum safeguards that should be set out in statute law in order to avoid abuses of power:

- the nature of the offences which may give rise to an interception order;
- a definition of the categories of people liable to have their telephones tapped;
- a limit on the duration of telephone tapping;
- the procedure to be followed for examining, using and storing the data obtained;
- the precautions to be taken when communicating the data to other parties; and
- the circumstances in which recordings may or must be erased or the tapes destroyed.

¹⁰ ECHR, 29 June 2006, Application no. 54934/00 (*Weber & Saravia v. Germany*).

Since 2006, a number of cases have been decided on the issue of surveillance, either in the context of post-crime or pre-crime measures, as well as measures taken by the intelligence services.¹¹ This includes both concrete interferences and legislation that would enable such interferences. As recounted above, the latter is not normally open to scrutiny by the Court, as it concerns an *abstract test* of the compatibility of domestic law against the Convention. The Court, however, can make an exception when applicants claim that the nature of the legislation or practice is such that they cannot know whether or not they have been a victim of state surveillance.

With the above analyses that closely follow the reasonings of the Court, the readers should have sufficient analytical instruments to study, for instance, the case of *Big Brother Watch and Others v. the United Kingdom* of 2018.¹² This case regards complaints about the compatibility with Article 8 ECHR of three discrete regimes of mass surveillance in the United Kingdom. First, the regime for the *bulk interception of communications* under section 8(4) of the Regulation of Investigatory Powers Act 2000 (RIPA); the UK–US intelligence sharing regime applied by the security service (MI5), the secret intelligence service (MI6), and the Government Communications Headquarters (GCHQ, which covers information and signals intelligence or ‘sigint’); and the regime for the *acquisition of communications data* under Chapter II of RIPA. The purpose of this work is not to provide an exhaustive overview of positive law in the realm of the right to privacy, but to provide computer scientists and students of computer science with a proper understanding of law as a scholarly discipline and a professional practice. In the end, the proof of the pudding will be in the eating. The reader is invited and encouraged to have their own tastings of legal texts, discovering the major impact of legal decision-making on potential violations of, for example, the right to privacy.

5.4 Privacy and Data Protection

Since the CFREU (or ‘the Charter’) has been in force (2009), the EU ‘has’ two fundamental rights regarding the processing of personal data:

¹¹ E.g. ECtHR, 1 July 2008, Application no. 58243/00 (*Liberty and Others v. the United Kingdom*); ECtHR, 18 May 2010, Application no. 26839/05 (*Kennedy v. the United Kingdom*); ECtHR, 4 December 2015, Application no. 47143/06 (*Roman Zakharov v. Russia*); ECtHR, 12 January, Application no. 2016 37138/14 (*Szabó and Vissy v. Hungary*); ECtHR, 19 June 2018, Application no. 35252/08 (*Centrum För Rättvisa v. Sweden*); ECtHR, 13 September 2018, Application nos. 58170/13, 62322/14 and 24960/15 (*Big Brother Watch and Others v. the United Kingdom*).

¹² ECtHR, 13 September 2018, Application nos. 58170/13, 62322/14 and 24960/15 (*Big Brother Watch and Others v. the United Kingdom*).

Article 7 Respect for private and family life

Everyone has the right to respect for his or her private and family life, home and communications.

Article 8 Protection of personal data

1. Everyone has the right to the protection of personal data concerning him or her.
2. Such data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right of access to data which has been collected concerning him or her, and the right to have it rectified.
3. Compliance with these rules shall be subject to control by an independent authority.

This is a new situation in the realm of human rights, because no other Constitution or Human Rights Treaty attributes a right to the protection of personal data.

Article 52 of the Charter clarifies the relationship between Article 7 of the Charter and Article 8 ECHR, which both refer to the right to privacy.

3. In so far as this Charter contains rights which correspond to rights guaranteed by the Convention for the Protection of Human Rights and Fundamental Freedoms, the meaning and scope of those rights shall be the same as those laid down by the said Convention. This provision shall not prevent Union law providing more extensive protection.

This stipulates that Article 7 CFREU cannot be interpreted as providing less protection compared to Article 8 ECHR, but may be interpreted as attributing additional protection. To the extent that Article 8 CFREU corresponds to Article 8 ECHR, it can—similarly—not be interpreted as providing less protection than Article 8 ECHR, but it may provide additional protection.

Before diving deep into the General Data Protection Regulation (GDPR) that provides more details, rules, and principles for the processing of personal data, we will first investigate how the fundamental right to data protection compares to the fundamental right to privacy.

5.4.1 Defaults: an opacity right and a transparency right

Some authors have argued that whereas, by default, the right to privacy is foremost an *opacity right*, data protection is foremost a *transparency right*. As an

opacity right, the right to privacy aims to safeguard a private sphere for individual citizens, where they can basically ward-off interference by others, most notably the state. This highlights the idea that privacy is a *liberty right*, a negative right that obligates others to refrain from interference with the good that is protected. As a transparency right, the right to data protection aims to ensure that whenever personal data is processed (which included collection, access, manipulation, and any other usage) such processing must be done in a transparent manner, in compliance with a set of conditions which should ensure fair and lawful processing.

Note that the opacity concerns the private sphere of an individual person, whereas the transparency concerns the state and other powerful actors when processing personal data. This accords with the core tenets of the Rule of Law, which hold that whereas government should be as transparent as possible, citizens should be shielded from intrusive transparency by the government.

Also, as discussed above, even though privacy is an opacity right that requires the state to refrain from interference (*negative freedom*), the right to privacy may, *nevertheless*, impose *positive obligations* on the state to enable individuals to exercise their right. Similarly, though data protection is a transparency right that should enable individuals as well as others to act on their personal data (*positive freedom*), while imposing a number of positive obligations on those who determine the purpose of processing, the right to data protection may, *nevertheless*, require that others abstain from processing personal data, thus imposing *negative obligations* on them.

5.4.2 Distinctive but overlapping rights: a Venn diagram

Though one may be tempted to see the right to data protection as a subset of the right to privacy, this is not correct. Within the context of the EU, the right to privacy entails both more and less than the right to data protection. We portray this in Figure 5.1 below.




Figure 5.1 Venn diagram of the fundamental rights to privacy and data protection

Whenever the processing of personal data constitutes an interference with the right to privacy, there is an overlap. The right to privacy, however, also concerns interference with bodily integrity, decisional privacy, privacy of the home, and correspondence when no processing of personal data is involved. This is where the right to privacy entails more than the right to data protection.

Similarly, the right to data protection also concerns the processing of personal data when there is no interference with the right to privacy, for instance, when one's personal data are processed on one's own request, for example, the processing of an address or banking details to deliver goods and charge one's account as a consequence of the sale of a book.

Note that if such data are subsequently used for other purposes, for example, to support the business model of a webshop by way of targeted advertising, privacy may be at stake. Whether or not this is the case also relates to the fact that the right to privacy, as discussed above, is primarily at stake in the vertical relationship between a government and its citizens, whereas the right to data protection seems to be applicable to all those who process personal data. This is certainly the case for data processing that falls under the scope of the GDPR.

5.4.3 Legal remedies in case of violation

The right to privacy can be invoked in a national court of law, for instance in the course of criminal or administrative proceedings. As discussed above, individual citizens have a right to present their claim to the ECtHR, which resides in Strasbourg, but this can only be done *after exhausting national remedies*. That means that if one fails to claim violation of Article 8 ECHR at the national level, or if one fails to appeal against a judgment that denies such a violation, the application to the ECtHR will be inadmissible. See Articles 34 and 35 ECHR:

Article 34 Individual applications

The Court may receive applications from any person, nongovernmental organisation or group of individuals claiming to be the victim of a violation by one of the High Contracting Parties of the rights set forth in the Convention or the Protocols thereto. The High Contracting Parties undertake not to hinder in any way the effective exercise of this right.

Article 35 Admissibility criteria

The Court may only deal with the matter after all domestic remedies have been exhausted, according to the generally recognised rules of international law, and within a period of six months from the date on which the final decision was taken.

Both the right to privacy and the right to data protection of the CFREU have direct application in the MSs of the EU. This means one can invoke them in a national court of law. If, however, a question is raised about the interpretation of the Charter, Article 267 of the Treaty on the Functioning of the EU (TFEU) stipulates that so-called ‘preliminary questions’ can, or must, be referred to the CJEU (which resides in Luxembourg):

The Court of Justice of the European Union shall have jurisdiction to give preliminary rulings concerning:

- (a) the interpretation of the Treaties;
- (b) the validity and interpretation of acts of the institutions, bodies, offices or agencies of the Union;

Where such a question is raised before any court or tribunal of a Member State, that court or tribunal may, if it considers that a decision on the question is necessary to enable it to give judgment, request the Court to give a ruling thereon.

Where any such question is raised in a case pending before a court or tribunal of a Member State against whose decisions there is no judicial remedy under national law, that court or tribunal shall bring the matter before the Court.

(...)

Clearly, both European Courts have an important role as to the national jurisdiction regarding human and fundamental rights. The case law of both Courts is a pivotal source of law that will remain central throughout this work.

5.5 Data Protection Law

The history of data protection law goes back to the 1970s, when various countries enacted legislation to ensure fair processing of personal information by the government. An early example was the US Privacy Act of 1974,¹³ which instigated a set of fair practices for dealing with personal information.

¹³ Privacy Act of 1974, 5 U.S.C. § 552a, see: <https://www.gpo.gov/fdsys/pkg/USCODE-2012-title5/pdf/USCODE-2012-title5-partI-chap5-subchapII-sec552a.pdf>.

In 1980, the global Organisation of Economic Co-operation and Development (OECD) issued the so-called 'Fair Information Principles' (FIPs), as part of the (non-binding) *Guidelines governing the protection of privacy and trans-border flows of personal data*:

Collection Limitation Principle

7. There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.

Data Quality Principle

8. Personal data should be relevant to the purposes for which they are to be used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.

Purpose Specification Principle

9. The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.

Use Limitation Principle

10. Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with Paragraph 9 except:
 - a) with the consent of the data subject; or
 - b) by the authority of law.

Security Safeguards Principle

11. Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.

Openness Principle

12. There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.

Individual Participation Principle

13. Individuals should have the right:
 - a) to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to them;

- b) to have communicated to them, data relating to them
 - i. within a reasonable time;
 - ii. at a charge, if any, that is not excessive;
 - iii. in a reasonable manner; and
 - iv. in a form that is readily intelligible to them;
- c) to be given reasons if a request made under subparagraphs (a) and (b) is denied, and to be able to challenge such denial; and
- d) to challenge data relating to them and, if the challenge is successful to have the data erased, rectified, completed or amended.

Accountability Principle

- 14. A data controller should be accountable for complying with measures which give effect to the principles stated above.

The version quoted has been taken from the updated Guidelines of 2013. The update does not concern the FIPs themselves, but aims to *strengthen world-wide enforcement and accountability*. With an eye to the increased scale of data processing and the new techniques for data analytics, the OECD recommends a risk-based approach that is proactive rather than reactive when it comes to the rights and freedoms of those affected by the processing of personal data.

Since 1980, many states have enacted data protection legislation, often following the FIPs. The EU Data Protection Directive (DPD) of 1995 was a prime example of a legally binding implementation of the OECD Guidelines. Since May 2018 the DPD has been succeeded by the GDPR. Just like the updated OECD Guidelines, the basic rules and principles that underlie the GDPR are largely the same as those of the DPD. The difference regards enforcement and various obligations to take a proactive approach to compliance. Again, a practical and effective reinforcement of the accountability principle is the most significant change.

5.5.1 EU and US data protection law

In the United States, data protection is part of the right to privacy (in Constitutional and tort law) and subject to sectorial legislation, notably with regard to finance, healthcare, special protection of children, and consumer protection. There is no general law on data protection, apart from the 1974 Privacy Act (which only applies to Federal Agencies). This means that the protection of personal data varies with the context of processing. In commercial contexts, much of the actual protection depends on the competences of the Federal Trade Commission (FTC), based on section 5 of the FTC Act:

- (1) Unfair methods of competition in or affecting commerce, and unfair or deceptive acts or practices in or affecting commerce, are hereby declared unlawful.
- (2) The Commission is hereby empowered and directed to prevent persons, partnerships, or corporations, [except certain specified financial and industrial sectors] from using unfair methods of competition in or affecting commerce and unfair or deceptive acts or practices in or affecting commerce.

Based on this, the FTC is tasked with protecting consumer privacy and data security in commercial contexts. The notion of a reasonable expectation of privacy is a core concept, because consumer trust is pivotal for a well-functioning market in ecommerce. The FTC deals with violations on a case-by-case basis, but also issues so-called ‘rulings’ if it believes specific types of violations are prevalent. Such ‘rulings’ basically declare how the FTC will use its Article 5 competence, thus encouraging companies to change their behaviour. The FTC is often qualified as ‘the regulator’ concerning informational privacy, due to its central role in US policy-making regarding data protection.

In the EU, the situation is altogether different, due to the general applicability of EU data protection law, which does not depend on whether a violation can be framed as ‘an unfair or deceptive act in or affecting commerce’. In the next subsection, we will provide an extensive discussion of the core content of EU data protection law.

One could say that whereas in the United States the processing of personal information is allowed unless it has been explicitly restricted, in the EU any processing of any personal data in any context is conditioned by a set of rules and principles that impose obligations on those who process data and attribute rights to those whose personal data is at stake.

5.5.2 EU data protection law

The GDPR is based on Article 16 TFEU, which reads:

1. Everyone has the right to the protection of personal data concerning them.
2. The European Parliament and the Council, acting in accordance with the ordinary legislative procedure, shall lay down the rules relating to the protection of individuals with regard to the processing of personal data by Union institutions, bodies, offices and agencies, and by the Member States when carrying out

activities which fall within the scope of Union law, and the rules relating to the free movement of such data. Compliance with these rules shall be subject to the control of independent authorities.

(...)

The GDPR protects the *fundamental right to data protection* as stipulated in Article 8 CFREU. However, the GDPR goes beyond this, explicitly aiming to protect *all the fundamental rights and freedoms* that are implicated by the processing of personal data. But this is not the only goal of the Regulation. At the same time, the Regulation aims to prevent that different levels of data protection within the jurisdiction of the MSs result in obstructions of the internal market. So, harmonization of protection to ensure a free flow of personal data is the second, equally important, goal of the GDPR:

Article 1 Subject-matter and objectives

1. This Regulation lays down rules relating to the protection of natural persons with regard to the processing of personal data and rules relating to the free movement of personal data.
2. This Regulation protects fundamental rights and freedoms of natural persons and in particular their right to the protection of personal data.
3. The free movement of personal data within the Union shall be neither restricted nor prohibited for reasons connected with the protection of natural persons with regard to the processing of personal data.

As discussed earlier (in section 4.3), the EU has developed from the European Economic Community (EEC), where the most important goal was the creation of an internal market, based on the ‘four freedoms’: free movement of capital, persons, goods, and services. As paragraph 3 of Article 1 GDPR clarifies, this Regulation involves ‘full harmonisation’, which means that MSs are not allowed to provide either less or more protection than what is offered in the Regulation (with the exception of explicitly formulated discretion). *Full harmonization* ensures the absence of obstructions of the internal market due to different requirements in terms of data protection. The fact that the GDPR is a regulation instead of a directive confirms the wish to eradicate such obstructions, thus hoping to boost data-driven business across national borders.

5.5.2.1 Sources of law regarding EU data protection law

So far, we have seen that the sources of law consist of legislation and treaties, case law, doctrine, customary law, and fundamental principles. In the case

of EU data protection law, we have the founding Treaties,¹⁴ the Charter, the GDPR, the Police Data Protection Directive (PDPD),¹⁵ the ePrivacy Directive (ePD),¹⁶ and a whole series of other Regulations and Directives that may at some point be relevant (but will not be discussed here). Next to this we have the case law of the CJEU regarding data protection issues, decisions and policies of the supervisory authorities in the MSs and the European Data Protection Supervisor, and we have doctrinal treatises and journal articles which analyse and discuss the legislation, the case law, and the underlying principles and practices.

In the case of EU data protection law, we have one more source of law, which has played an important role in the interpretation of the former DPD: the Opinions and Guidelines of the independent Article 29 Working Party (Art. 29 WP). This was the advisory body (instituted by Article 29 DPD) that produced a great number of highly relevant interpretations of EU data protection law, which continue to function as an important source of law. Though its output was not binding, it has persuasive authority based on the experience and expertise of its members (the data protection supervisors of the MSs) and based on its official task, which was to advise on proper implementation of EU data protection law. Most of the Opinions, Guidelines and Recommendations of the Art. 29 WP are equally relevant under the GDPR, as the core principles and concepts have not changed.

The Art. 29 WP has been replaced, under the GDPR, with the independent European Data Protection Board (EDPB),¹⁷ instituted in Articles 68–76 GDPR, again consisting of the supervisory authorities of all the MSs of the EU, again tasked with advising on the correct interpretation of the EU data protection law. The EDPB has further tasks in contributing to a harmonized approach of the national supervisors, throughout the Union.

¹⁴ European Union, Treaty on European Union (Consolidated Version), Treaty of Maastricht, 7 February 1992, Official Journal of the European Communities C 325/5; 24 December 2002, available at: <http://www.refworld.org/docid/3ae6b39218.html>, for the TFEU see above footnote 8 in Chapter 4.

¹⁵ Directive (EU) 2016/680 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data, and repealing Council Framework Decision 2008/977/JHA.

¹⁶ Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications).

¹⁷ <https://edpb.europa.eu>, its Opinions and Guidelines can be found at: https://edpb.europa.eu/our-work-tools/general-guidance/gdpr-guidelines-recommendations-best-practices_en.

5.5.2.2 Material and territorial scope

The material scope of the GDPR is limited to ‘the processing of personal data’ (Article 2.1). The definition of ‘processing’, however, is very broad, as Article 4(2) reads:

‘processing’ means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction

The GDPR does not apply to the processing of personal data within the context of a household and it does not apply to processing of personal data in the context of the prevention and prosecution of crime and threats to public security.¹⁸ The *household exception* will usually exempt the users of social networks, but not the providers (see section 5.5.2.4). With regard to the prevention and prosecution of crime, the PDPD is in force, based on Article 39 of the Treaty of the European Union (TEU).¹⁹ Since the EU has no competence regarding public security (intelligence services), there is no EU legislation as to the processing of personal data in the context of threats to public security. Note that the ECHR does apply to issues of public security, so insofar as privacy is infringed, measures can be tested against Article 8 ECHR (see section 5.3.5, notably the cases of *Klass* and *Weber & Saravia*).

Next to the exemptions of Article 2, Article 33 states that MSs may enact legislation to restrict the applicability of specific GDPR provisions, if they regard measures that are necessary in a democratic society, targeting a limited set of goals, such as national security, defence, public security, the prevention, investigation, detection, and prosecution of criminal offences, or of breaches of ethics for regulated professions, an important object of general public interest of a MS or of the EU, including monetary, budgetary, and taxation matters. Note that though restrictions based on these goals are allowed if they pass the proportionality test (‘necessary in a democratic society’ clearly refers to Article 8.2 ECHR), they also require a basis in law.

¹⁸ Art. 2 GDPR.

¹⁹ Directive (EU) 2016/680 on the protection of individuals with regard to the processing of personal data by competent authorities for the purposes of prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and the free movement of such data.

Any such restrictions are only valid insofar as they respect *the essence of the fundamental rights and freedoms*.

The territorial scope of the GDPR is defined in Article 3:

1. This Regulation applies to the processing of personal data in the context of the activities of an establishment of a controller or a processor in the Union, regardless of whether the processing takes place in the Union or not.

Bear in mind that if a tech company has an establishment in the EU, the GDPR applies to the processing of personal data, even if the processing takes place elsewhere. At some point a tech company relocated its headquarters from Ireland to the United States, because otherwise data subjects in countries outside the EU could appeal to the Irish data protection supervisor under the GDPR.

2. This Regulation applies to the processing of personal data of data subjects who are in the Union by a controller or processor not established in the Union, where the processing activities are related to:
 - (a) the offering of goods or services, irrespective of whether a payment of the data subject is required, to such data subjects in the Union; or
 - (b) the monitoring of their behaviour as far as their behaviour takes place within the Union.

Here we see that if a company decides to offer goods or services (whether or not they are free) that involve the processing of personal data of data subjects in the Union, or monitor their behaviour in the Union, the GDPR applies, irrespective of whether the company is established in the Union. Consider that this jurisdiction is limited to data subjects who are in the Union; it does not apply to EU citizens outside the Union, though it does apply to non-EU citizens when they are in the Union.

5.5.2.3 Personal data and data subject

Article 4(1) GDPR clarifies that:

‘personal data’ means:

- any information
- relating to
- an identified or
- identifiable natural person (‘data subject’)

where ‘an identifiable person’ is defined as

- one who can be identified,
- directly or indirectly,
- an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity;

Many authors have pointed out that this entails a very broad view of ‘personal data’, potentially bringing nearly any data under the heading of personal data. This is especially the case as the combination of increased availability and increased searchability and linkability of massive amounts of data will enable identification and re-identification of data that would previously not have been considered personal data. At some point, data about the weather, about room temperature, about the arrival of a train may become personal data, when it can be related to a person that can be singled out. Recital 26 reads:

To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly.

The criterion to determine whether data is personal, that is ‘identifiable’, is that it is ‘reasonably likely’ that a person can, for example, be singled out. The recital continues:

To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments.

Here we see that the ‘reasonably likely’ criterion should be understood as an *objective criterion*, taking into account the costs, the time, and effort, and the available technical means at the time of processing.

In the case of *Breyer v. Germany*,²⁰ the CJEU decided that even a dynamic IP address may qualify as a personal data, depending on whether the link with a

²⁰ CJEU, 19 October 2016, Case 582/14 (*Patrick Breyer v. Germany*).

specific person can be made. The case concerned government websites that processed dynamic IP addresses, keeping them longer than was necessary for providing access to the sites. What made this case special is that the ability to link the IP address to a specific person was not in the hands of the operators of the government website but in the hand of internet service providers (ISPs). The CJEU found that because ISPs could be ordered by a court to provide information about the user of a dynamic IP addresses, this IP address should not be considered anonymous.

So, *personal data* is any data that relates to an identifiable natural person (excluding legal persons such as corporations), and a *data subject* is the identifiable natural person to whom the data relate.

The material scope of the GDPR regards (as discussed in section 5.5.2.2) the processing of personal data. This implies that to avoid applicability of the GDPR, one could ‘simply’ anonymize previously personal data. There are two caveats here. First, *anonymization* is itself a form of processing, and thereby requires a valid legal ground (see section 5.5.2.5). Second, anonymization is not easy, because the risk of re-identification easily turns ‘anonymous’ data into identifiable and thus personal data. In practice, anonymization will often remove so much information from the data that it is no longer relevant for the purpose of processing. To better understand the difference between personal and anonymized data, we can best check the definition of ‘pseudonymization’ of Article 4(5):

- the processing of personal data in such a manner that
- the personal data can no longer be attributed to a specific data subject
- without the use of additional information,
- provided that such additional information is kept separately and
- is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person;

First, we see that *pseudonymous* data is defined as a subset of personal data. Second, it is defined as data from which any identifying information has been removed and stored separately, subject to technical and organizational measures that resist re-identification.

Pseudonymization is a way to comply with data protection law (by ways of data minimization), not a way to avoid applicability. With regard to encryption, key management that enables a party other than the data subject to decrypt, will mostly qualify

as pseudonymization, not as anonymization. Bear in mind that the definition of pseudonymization in the GDPR defines the condition of the relevant legal effect, irrespective of other how other disciplines define pseudonymization.

5.5.2.4 Data controller and data processor

Article 4(7) GDPR defines ‘controller’ as:

- the natural or legal person, public authority, agency or any other body
- which alone or jointly with others
- determines the purposes and means of the processing of personal data;

The definition of ‘data controller’ is crucial, because the ‘data controller’ is both accountable and liable for compliance with all the obligations of the GDPR, including obligations to implement a proactive approach to potential risks to the fundamental rights and freedoms of data subjects. The ‘data controller’ is basically defined as whoever determines the purpose of processing, whereby the CJEU checks who determines such purpose in practice, not merely on paper. The ‘data controller’ also determines the means of processing, but this can be outsourced to a data processor, defined as (Article 4(8)):

- a natural or legal person, public authority, agency or any other body
- which processes personal data on behalf of the controller;

Here, we clearly see that the data controller remains accountable for the choice of the means of processing, even if that choice is made by a processor. When the landmark case on the so-called ‘right to be forgotten’ was decided in 2014 (*Google Spain v. Costeja*),²¹ one of the most important issues was whether Google should be qualified as a data controller or a data processor. Google had argued that its search engine has no other function than to provide its users with automatically generated search results, thereby claiming that it is the user, not the service provider who determines the purpose of the processing. Google argued that its search engine is merely a choice of means (the PageRank algorithm) employed in the service of users that decide the purpose of the search. The highest adviser of the CJEU (the Court), holding the office of Advocate General (AG), who is required to provide a so-called ‘Opinion’ (advise) to the Court, had taken the position that in the case of a search engine, the service provider is indeed a data processor, not the controller. Surprisingly,

²¹ CJEU, 13 May 2014, C-131/12 (*Google Spain v. Costeja*). See also EDPB (formerly Art. 29 WP), Opinion 1/2010 on the concepts of ‘controller’ and ‘processor’.

the Court (that is not bound by the Opinion of the AG), took another position, based on the fact that Google Spain (the subsidiary that sells advertising space on the search engine's pages directed to Spanish users) has its own business model and thereby determines the purpose of processing. If the Court had not qualified Google Spain as a data controller, it could never have required Google to de-list the news item that Costeja wished to have erased.

Another pivotal case of 2018 concerned the fanpage of *Wirtschaftsakademie*,²² used to provide services in the realm of education. The fanpage was hosted on Facebook, which enabled the operator to obtain anonymous statistical details on website visitors via the 'Facebook Insights' function, which Facebook offers free of charge under non-negotiable conditions. The CJEU decided that the operator of the fanpage was a *joint controller*, together with Facebook, as the statistics were obtained by processing cookies placed on the terminal equipment of the visitors. Since the purpose of the processing of such cookies is co-decided by the fanpage operator, even though it has no control over the data processing and was not given access to the data, they are jointly responsible for the necessary processing of personal data. Under the GDPR this would be based on Article 26, which reads:

Where two or more controllers jointly determine the purposes and means of processing, they shall be joint controllers. They shall in a transparent manner determine their respective responsibilities for compliance with the obligations under this Regulation, in particular as regards the exercising of the rights of the data subject and their respective duties to provide the information (...).

As in the case of *Google Spain v. Costeja*, where the AG argued that Google was merely a processor, acting on request of the users of the search engine, one could argue that, in this case, Facebook is acting as a processor for the fanpage operator who wishes to obtain the statistics. In line with the Court in *Google Spain v. Costeja*, the Court decided that Facebook is the controller, not the processor. In this case, however, the fanpage operator—other than the users of a search engine—is considered a joint controller.²³

²² CJEU, 5 June 2018, C-210/16 (*Unabhängiges Landeszentrum für Datenschutz Schleswig-Holstein v. Wirtschaftsakademie Schleswig-Holstein GmbH*).

²³ See also CJEU, 29 July 2019, Case C-40/17 (*Fashion ID*), where the Court ruled that 'The operator of a website, such as Fashion ID GmbH & Co. KG, that embeds on that website a social plugin causing the browser of a visitor to that website to request content from the provider of that plugin and, to that end, to transmit to that provider personal data of the visitor can be considered to be a controller, within the meaning of Article 2(d) of Directive 95/46. That liability is, however, limited to the operation or set of operations involving the processing of personal data in respect of which it actually determines the purposes and means, that is to say, the collection and disclosure by transmission of the data at issue.'

5.5.2.5 Legal ground for lawful processing of personal data

The processing of personal data is only allowed on the basis of one of six legal grounds. Please take note of the fact that consent is just one of those legal grounds and not necessarily the most obvious. Article 6 GDPR reads:

- a) the data subject has given consent to the processing of his or her personal data for one or more specific purposes;

Under the DPD ‘for one or more specific purposes’ was not explicitly mentioned, though it was obvious from requirements detailed elsewhere in the DPD. Under the GDPR, it is explicitly clear that consent is only valid if the purpose has been specified. As Article 5 stipulates that data may only be processed if necessary for the specified purpose, this means that consent can only concern the processing of personal data that is necessary for the purpose that was communicated. All the other grounds stipulate that the processing must be necessary in relation to the ground.

Valid consent will be further discussed in a dedicated section (section 5.5.2.7).

- b) processing is necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract;

This entails that once the contract has been concluded and performed and the data is no longer necessary (goods or service delivered, invoice paid), it may no longer be processed on this ground. Further processing will require another ground, for example, consent (for another purpose).

- c) processing is necessary for compliance with a legal obligation to which the controller is subject;

Much processing is mandatory due to legal obligations, such as processing by the tax authority, social security agency, land registry, or by commercial enterprise that must, for example, comply with employment, social security, and tax legislation. Article 6.3 stipulates that this processing must be based on MS or Union law, must contain the specific purpose(s) of processing, and must have relevant limitations and safeguards.

- d) processing is necessary in order to protect the vital interests of the data subject or of another natural person;

This ground must be understood as concerning life-threatening situations, where, for example, medical data must be processed to save someone's life.

- e) processing is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller;

This ground is comparable to the c-ground, but here there may not be a legal obligation but a legal competence or task that requires the processing of personal data. We can think of processing by various types of government agencies that provide support to those in need, or need to collect information on energy usage to develop policies on the reduction of energy consumption. Note that to the extent that such information can rely on aggregated or otherwise anonymized data, the processing of personal data is not necessary and cannot be based on this ground. Article 6.3 stipulates that this processing must be based on MS or Union law, must contain the specific purpose(s) of processing, and must have relevant limitations and safeguards.

- f) processing is necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.

Point (f) of the first subparagraph shall not apply to processing carried out by public authorities in the performance of their tasks.

The f-ground is important for processing carried out by the commercial sector, including financial institutions, social networks, and search engines, and we may expect that added value service providers in the context of smart homes, smart grids, and connected cars will base the processing of data that is not necessary for the primary process (which will often be based on contract) on the f-ground. As the *economic interests of a business*, including its competitive edge and innovative potential, often depend on advertising revenue and/or the sale of personal data or inferred profiles, the f-ground is a tempting basis insofar as other grounds do not apply.

However, the f-ground requires a *balancing test*. As one can imagine, the business interests of a company cannot, by default, overrule the interests or fundamental rights and freedoms of data subjects whose behavioural data are used to generate income (thus, e.g. enabling the so-called 'free services' of social networks and search engines).

This has two consequences:

1. The controller has to assess (before initiating the processing) whether its economic interests in processing personal data that are not necessary to provide a requested service, can overrule the interests and rights of those whose data are used for micro-targeting or other ways of monetizing the data.
2. The data subject can object to the processing based on their particular situation, in which case the controller must stop processing unless it can find compelling grounds to override the interests, rights, and freedoms of the data subject. The right to object is based on Article 21 GDPR and also concerns the e-ground.

The balancing test required of the controller, entails the following considerations:²⁴

the nature and source of the legitimate interest and whether the data processing is necessary for the exercise of a fundamental right, is otherwise in the public interest, or benefits from recognition in the community concerned;

the impact on the data subject and their reasonable expectations about what will happen to their data, as well as the nature of the data and how they are processed;

additional safeguards which could limit undue impact on the data subject, such as data minimisation, privacy-enhancing technologies; increased transparency, general and unconditional right to opt-out, and data portability.

The f-ground is often used to legitimize the advertising business model of free services. For instance, in the *Google Spain v. Costeja* case discussed above, the CJEU concluded that Google was processing personal data based on its legitimate business interests. In this particular case, the CJEU considered two types of legitimate interests that might overrule Costeja's interest in having a particular search result de-referenced. First, the *interest of the controller*, second the *interests of third parties*, namely the users of the search engine.

First, the Courts looks into the economic interest of Google Spain in sustaining its business model, because the right to erase and the right to object

²⁴ EDPB (formerly Art. 29 WP), Opinion 06/2014 on the notion of legitimate interests of the data controller under Article 7 of Directive 95/46/EC, WP217, at 3.

that Costeja invoked would involve costs on the side of Google (especially because many others may similarly submit requests to de-reference). The CJEU found that:

81 In the light of the potential seriousness of that interference, it is clear that it cannot be justified by merely the economic interest which the operator of such an engine has in that processing. (...)

The seriousness of the interference, in this case, was argued in considerations 37 and 38:

37 (...) the organisation and aggregation of information published on the internet that are effected by search engines with the aim of facilitating their users' access to that information may, when users carry out their search on the basis of an individual's name, result in them obtaining through the list of results a structured overview of the information relating to that individual that can be found on the internet enabling them to establish a more or less detailed profile of the data subject.

38 Inasmuch as the activity of a search engine is therefore liable to affect significantly, and additionally compared with that of the publishers of websites, the fundamental rights to privacy and to the protection of personal data, (...).

Second, the Court considered the legitimate interests of users of the search engine in having access to the search result that may be de-referenced:

81 (...) However, inasmuch as the removal of links from the list of results could, depending on the information at issue, have effects upon the legitimate interest of internet users potentially interested in having access to that information, in situations such as that at issue in the main proceedings a fair balance should be sought in particular between that interest and the data subject's fundamental rights under Articles 7 and 8 of the Charter. Whilst it is true that the data subject's rights protected by those articles also override, as a general rule, that interest of internet users, that balance may however depend, in specific cases, on the nature of the information in question and its sensitivity for the data subject's private life and on the interest of the public in having that information, an interest which may vary, in particular, according to the role played by the data subject in public life.

Here we see a clash between the freedom of information of search engine users and the right to data protection of the data subject, which requires some subtle

balancing. Note, however, that the Court is not discussing the removal of content from the internet, but the de-referencing of a search result that links to such content.

The EDPB considers that in principle a controller must make up its mind which legal basis justifies a particular type of processing operation; controllers cannot, for instance, process personal data based on consent and then shift to the legitimate interest after the data subject withdraws their consent.²⁵ The EDPB also refers to Articles 13.1 and 14.1 that require the controller to provide information about the purpose(s) and the legal basis for its processing operations, meaning that this should be clarified from the start.²⁶

5.5.2.6 Principles of lawful, fair, and transparent processing

Next to, and thus on top of, having a legal ground, Article 5 GDPR stipulates a set of rules under the heading of ‘Principles relating to the processing of personal data.’ Though the use of the term ‘principles’ could suggest that these are just some underlying assumptions, they are in fact rules that must be complied with. We will follow the wording of the article, discussing each paragraph along the way (the principles in bold are part of the article, emphasis is mine):

- (a) processed lawfully, fairly and in a transparent manner in relation to the data subject (**‘lawfulness, fairness and transparency’**);

Though one may think that *lawfulness* merely refers to Article 6, which contains the legal basis, the term ‘lawfulness’ also refers to the bigger picture of the *rule of law*, as with the requirement that infringements of the right to privacy under Article 8 ECHR must be ‘in accordance with the law’. This means that a mere basis in law is not enough and must be understood in qualitative terms to include respect for *legitimate expectations*, *independent oversight*, and other *checks and balances* to ensure that the legal basis of Article 6 is valid (see also Article 6.3). Similarly, fairness refers to various balancing and proportionality tests, taking note of the relevant interests and fundamental rights that are at stake. *Transparency* is further detailed in Articles 13, 14, and 15 GDPR.

²⁵ EDPB (formerly Art. 29 WP), 10 April 2018, Guidelines on consent under Regulation 2016/679, WP259.rev.1, 23.

²⁶ EDPB (formerly Art. 29 Working Party), 10 April 2018, Guidelines on transparency under Regulation 2016/679, WP260.rev.1, 33, 35.

(b) collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes; further processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes shall, in accordance with Article 89(1), not be considered to be incompatible with the initial purposes (**‘purpose limitation’**);

Purpose limitation is one of the most important principles of EU data protection law. The idea that data must be collected and processed for one or more legitimate purposes that have been made explicit and are sufficiently specified pervades the regulation, while at the same time qualifying whoever—de facto—determines such purpose(s) as the responsible, accountable, and liable entity (the data controller). *Purpose is, in a way, the vanishing point of the architecture of EU data protection law.*

Further processing for another purpose is allowed if the purpose is *not incompatible* with the initial purpose, as communicated to the data subject. To determine whether the new purpose is compatible, Article 6(4) provides the following indications: any link between the old and the new purpose, the context of collection and the relationship between controller and subject, the nature and sensitivity of the data, the potential consequences of further processing for the data subject, and the existence of appropriate safeguards, such as encryption or pseudonymization. In case of consent for the new purpose or a legal obligation that involves the new purpose, processing is based on the new ground and cannot be based on processing for a compatible purpose.

Secondary usage (further processing) for *scientific or statistical research or archiving in the public interest* is considered compatible by default. The GDPR contains an extensive exception for such processing in Article 89, with further exceptions for medical research in, for example, Article 9.2(h). Recital 33 furthermore indicates that ‘[i]t is often not possible to fully identify the purpose of personal data processing for scientific research purposes at the time of data collection. Therefore, data subjects should be allowed to give their consent to certain areas of scientific research when in keeping with recognised ethical standards for scientific research. Data subjects should have the opportunity to give their consent only to certain areas of research or parts of research projects to the extent allowed by the intended purpose.’

(c) adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed (**‘data minimisation’**);

Data minimization is another core principle, that also underlies the principles of purpose limitation and storage limitation. In the DPD, this ground was articulated as ‘adequate, relevant and not excessive’, whereas now the criterion is ‘adequate, relevant and limited to what is necessary’. This is a further restriction, moving towards strict proportionality and subsidiarity, thereby also relating to the requirement to pseudonymize or anonymize the data as soon as possible. This principle links consent to necessity, as observed above. It also connects with the right to request erasure if processing is irrelevant for the given purpose.

(d) accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay (**‘accuracy’**);

Here the principle of *accuracy* is formulated as a legal obligation of the data controller, but this connects with the rights of erasure and rectification in the case that data are inaccurate.

(e) kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed; personal data may be stored for longer periods insofar as the personal data will be processed solely for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes in accordance with Article 89(1) subject to implementation of the appropriate technical and organisational measures required by this Regulation in order to safeguard the rights and freedoms of the data subject (**‘storage limitation’**);

Storage limitation basically requires that controllers engage in lifecycle management of the personal data they process, removing them, for example, when the purpose is exhausted and processing is no longer relevant. The exception for scientific research and archiving, mentioned above, requires appropriate technical and organizational safeguards, taking into account the rights and freedoms of the data subject, which will vary depending on, for example, the nature of the data.

(f) processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures (**‘integrity and confidentiality’**).

This principle connects with the requirement of *security by design* of Article 32, and the legal obligation for controllers to notify supervisory authorities and data subjects of data breaches (Articles 33, 34).

2. The controller shall be responsible for, and be able to demonstrate compliance with, paragraph 1 (**'accountability'**).

The *accountability principle* addresses the data controller as the focal point of responsibility, accountability, and liability regarding compliance with the principles that pervade the GDPR. Accountability is further detailed in Article 30 that requires the controller to demonstrate and document compliance, while liability is further detailed in Articles 79–83 about enforcement (including both administrative law fines and prohibitions, and private law compensation and injunctive relief). The roles and responsibilities of the controller (including joint controllers) and processor are further specified in Articles 24, 26, and 28.

5.5.2.7 Valid consent

Other than the DPD, the GDPR contains a separate article on consent. Article 7 declares, under the heading of 'Conditions for Consent':

1. Where processing is based on consent, the controller shall be able to demonstrate that the data subject has consented to processing of his or her personal data.

This concerns the burden of proof.

2. If the data subject's consent is given in the context of a written declaration which also concerns other matters, the request for consent shall be presented in a manner which is clearly distinguishable from the other matters, in an intelligible and easily accessible form, using clear and plain language. Any part of such a declaration which constitutes an infringement of this Regulation shall not be binding.

Note that consent may not be hidden in complicated wordy privacy policies, and must be 'easily accessible' as to its form (think of the user interface), 'using clear and plain language'. If consent is part of an elaborate and incomprehensible Terms of Service that basically contains an implicit consent, such consent is not valid.

3. The data subject shall have the right to withdraw his or her consent at any time. The withdrawal of consent shall not affect the lawfulness of processing based on

consent before its withdrawal. Prior to giving consent, the data subject shall be informed thereof. It shall be as easy to withdraw as to give consent.

This means that if consent is given by ticking a box, it must be as easy to untick the box. If one has to explore every nook and corner of a website to figure out how to withdraw consent, the consent is not valid.

4. When assessing whether consent is freely given, utmost account shall be taken of whether, inter alia, the performance of a contract, including the provision of a service, is conditional on consent to the processing of personal data that is not necessary for the performance of that contract.

To better understand what this means, we can use recital 43:

In order to ensure that consent is freely given, consent should not provide a valid legal ground for the processing of personal data in a specific case where there is a clear imbalance between the data subject and the controller, in particular where the controller is a public authority and it is therefore unlikely that consent was freely given in all the circumstances of that specific situation. Consent is presumed not to be freely given if it does not allow separate consent to be given to different personal data processing operations despite it being appropriate in the individual case, or if the performance of a contract, including the provision of a service, is dependent on the consent despite such consent not being necessary for such performance.

This seems to indicate that attempts to force consumers to choose between accessing a service and refusing consent for additional processing are unlawful and that such consent is not valid. Additional processing refers to the processing of data that is not necessary for the provision of the service, or further processing of data after the purpose has been exhausted. One could guess that Article 7.4, read through the lens of recital 43, is the end of a specific type of business model that puts the site on black if consent for unnecessary processing is not given.

Note that the legal ground must be communicated to the data subject *when the processing commences* (if data is collected from the data subjects, cf. Article 13), or within a reasonable time, at the latest within one month after obtaining the data (if data has not been obtained from the data subjects, cf. Article 14). Controllers cannot require consent and—after finding that the consent is not valid—claim that the processing is based on its legitimate interest; due to the

inherent logic of the different grounds, controllers cannot claim to base the same processing operations on different grounds.

5.5.2.8 Special categories of data

Article 9 defines a set of data as requiring special treatment. These data are often called ‘sensitive data’ and are defined as: ‘data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person’s sex life or sexual orientation’.

By default, the processing of such data is prohibited. Strictly defined exceptions apply, notably based on explicit consent; specific rights and obligations in the field of employment and social; the vital interests of the data subject or of another natural person; or with regard to processing in the context of not-for-profit bodies with a political, philosophical, religious, or trade union aim; processing of personal data which are manifestly made public by the data subject; processing necessary for legal claims, substantial public interest, preventive or occupational medicine, assessment of the working capacity of the employee, medical diagnosis, the provision of health or social care or treatment or the management of health or social care systems and services, for public health, for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes.

On top of that, Article 10 restricts the ‘Processing of personal data relating to criminal convictions and offences’.

Processing of personal data relating to criminal convictions and offences or related security measures based on Article 6(1) shall be carried out only under the control of official authority or when the processing is authorised by Union or Member State law providing for appropriate safeguards for the rights and freedoms of data subjects. Any comprehensive register of criminal convictions shall be kept only under the control of official authority.

Articles 9 and 10 demonstrate that data protection is not just about the right to privacy, but also entails protection against discrimination on prohibited grounds.

This is particularly relevant if inferences are made based on machine learning or other techniques to infer patterns from big data, because such inferences

may include sensitive data. Social networks, advertising intermediaries, or criminal justice authorities may infer racial or ethnic origin, political opinion, or sexual preferences, which inferences may then be applied to identifiable persons that match the profile. Such inferencing may be inadvertent, but nevertheless result in decisions based on such inferences, for instance parole decisions based on a correlation between race and recidivism. In Chapters 10 and 11 we will return to this point when discussing machine learning and profiling, including an analysis of GDPR provisions on profiling and automated decision-making based on profiling.

5.5.2.9 Data protection by design and default (DPbDD)

In Chapter 1, notably in section 1.4, we have identified the text-driven nature of modern law, in contrast with the orality of prior normative orderings. The rise of data- and code-driven ICIs confronts the text-driven nature of the law with a number of problems. Merely writing down and enacting legal norms may not work if the defaults of the technical and organizational architecture of the onlife world generate a contradictory normativity, which renders compliance with legal norms difficult if not impossible. In other words, the technical architecture may present its users and inhabitants with a *choice architecture* that limits their understanding of the backend systems of the social networks they use, of their smart homes, connected cars, and more.

Article 25 GDPR requires that data controllers design the data processing operations in compliance with data protection law. Data protection by design (DPbD) may sound like Privacy by Design (PbD). However, the latter is based on an ethical duty, not necessarily on a legal obligation; PbD reflects the choice of a controller to respect the privacy of their users by way of a privacy-friendly design. Also, as privacy is not equivalent with data protection, PbD cannot be equated with DPbD, even though in practice the terminology is often used interchangeably.

DPbD is a new legal obligation (no such obligation applied under the DPD). In case of non-compliance the legal effect is *liability for damages* (private law liability, Article 82), *unlawful processing* (administrative fines, Article 83), or *injunctive relief* (private law injunction to stop unlawful processing with penalty payments for every day of non-compliance, Article 79).

Under the heading of ‘data protection by design and default’, Article 25 stipulates:

1. Taking into account the state of the art, the cost of implementation and the nature, scope, context and purposes of processing as well as the risks of varying

likelihood and severity for rights and freedoms of natural persons posed by the processing, the controller shall, both at the time of the determination of the means for processing and at the time of the processing itself, implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects.

Paragraph 1 describes ‘data protection by design’ as a set of technical and organisational measures that embed core data protection principles into the design of the data processing architecture.

This should mitigate potential risks for the rights and freedoms of natural persons. The latter demonstrates the risk-based approach of the GDPR, which requires that controllers take a *proactive approach* when developing their computational backend systems. Note that Article 25 does not speak of the risks for rights and freedoms of data subjects, but of natural persons. This includes processing operations that impact other individuals, for instance when inferencing behavioural correlations that enable the influencing, exclusion, or other types of targeting of *others than the data subject*. Relevant design measures are, for instance pseudonymization, but one can also think of user-friendly interfaces to enable easy withdrawal of consent (Article 7.3) or subject access requests (SARs) (based on Article 15.3). Both the withdrawal of consent and SARs will involve computational architectures in the backend systems that effectively halt the processing of data for which consent has been withdrawn, or provide the data that are being processed (where Article 15.3 stipulates that if a SAR is made via electronic means, the data shall be provided in a commonly used electronic format).

The legal obligation to implement DPbDD is not obvious and may result in a major upheaval of backend systems, involving substantial costs. Depending on the risks of abstaining from such measures for the rights and freedoms of natural persons, such costs will become part of the relevant business model.

2. The controller shall implement appropriate technical and organisational measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are processed. That obligation applies to the

amount of personal data collected, the extent of their processing, the period of their storage and their accessibility. In particular, such measures shall ensure that by default personal data are not made accessible without the individual's intervention to an indefinite number of natural persons.

Paragraph 2 describes 'data protection by default', which is DPbD with regard to data minimization.

It demands that the architecture is constructed in such a way that no additional processing takes place, beyond what is necessary for the specific purpose of the relevant processing operations. Again, compliance with this legal obligation will result in major reconfigurations of current backend systems, involving, for example, effective lifecycle management of personal data (including pseudonymization, anonymization, and deletion of data).

The third paragraph declares that an approved certification mechanism may contribute to demonstration of compliance with DPbDD.

5.5.2.10 Data protection impact assessment

DPbDD is closely related to another new compliance mechanism, the data protection impact assessment (DPIA), again exhibiting the risk-based, proactive approach that is favoured under the GDPR. Basically, controllers are obligated to assess potential violations of the GDPR when initiating new data-driven technologies. Article 35 reads under the heading of 'data protection impact assessment' that:

1. Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data. A single assessment may address a set of similar processing operations that present similar high risks.

The criterion that decides whether a controller must conduct a DPIA is that foreseen processing operations are 'likely to result in a high risk to the rights and freedoms of natural persons'. Again, these risks are not restricted to data subjects, but extend to all natural persons. The assessment investigates the potential impact of envisaged processing operations, which assumes that these are indeed foreseen and mapped against impact on fundamental rights.

2. The controller shall seek the advice of the data protection officer, where designated, when carrying out a data protection impact assessment.

Articles 37–39 detail which types of controller must appoint a data protection officer (DPO), under what conditions (e.g. safeguards for independence) and with what tasks. One of the tasks of the DPO is to advise on the DPIA.

3. A data protection impact assessment referred to in paragraph 1 shall in particular be required in the case of:
- (a) a systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person;
 - (b) processing on a large scale of special categories of data referred to in Article 9(1), or of personal data relating to criminal convictions and offences referred to in Article 10; or
 - (c) a systematic monitoring of a publicly accessible area on a large scale.

Paragraph 3 sums up when a DPIA is mandatory, thus also giving an indication of what types of processing operations are considered *high-risk*.

Paragraphs 4–6 stipulate that supervisory authorities shall publish a further list of the kind of processing operations where a DPIA is mandatory, and may publish a list of processing operations where a DPIA is not mandatory. Both lists will be shared with the EDPB (which has an important advisory function as to the interpretation of the GDPR, and is further defined in Articles 68–76 GDPR).

7. The assessment shall contain at least:
- (a) a systematic description of the envisaged processing operations and the purposes of the processing, including, where applicable, the legitimate interest pursued by the controller;
 - (b) an assessment of the necessity and proportionality of the processing operations in relation to the purposes;
 - (c) an assessment of the risks to the rights and freedoms of data subjects referred to in paragraph 1; and
 - (d) the measures envisaged to address the risks, including safeguards, security measures and mechanisms to ensure the protection of personal data and to demonstrate compliance with this Regulation taking into account the rights and legitimate interests of data subjects and other persons concerned.

Paragraph 7 provides a first indication of a template for the DPIA. The listing has a high level of abstraction, thus enabling adequate concretization, depending on the types of processing operations, the context of processing, the nature of the data, and so forth. Under (d) we recognize a reference to DPbD, whose purpose is to mitigate risks to the rights and freedoms of natural persons.

Paragraph 8 states that approved codes of conduct (Article 40 GDPR) will be taken into account when assessing the impact.

9. Where appropriate, the controller shall seek the views of data subjects or their representatives on the intended processing, without prejudice to the protection of commercial or public interests or the security of processing operations.

Paragraph 9 emphasizes the need to involve those who will suffer the consequences of the intended processing, both on the side of data subjects and on the side of the controller. In earlier versions of the GDPR, the need to involve data subjects in the assessment was articulated more forcefully. One can imagine that a robust architecture will fare well based on input from those who will be effectively affected.

Paragraph 10 provides an exception for processing based on a legal obligation or a public task or authority (Article 6.1 under (c) and (e)), whenever the enactment of such an obligation has been preceded by a general DPIA on account of the legislator.

11. Where necessary, the controller shall carry out a review to assess if processing is performed in accordance with the data protection impact assessment at least when there is a change of the risk represented by processing operations.

In an innovative environment, where agile computing strategies lead to iterant changes in processing operations, the DPIA is best seen as a persistent and dynamic process that continuously monitors both the foreseeable risks and the appropriate mitigating measures.

5.5.2.11 Compliance and enforcement

The GDPR reinforces the accountability principle by initiating new legal obligations to further compliance, notably the obligation to implement DPbDD and to conduct a DPIA. Apart from those, other legal obligations require technical and organizational compliance measures, such as easy withdrawal of

consent (Article 7.3), provision of access by way of an electronic file (Article 15.3), obligations to employ pseudonymization (Articles 6.4(e), 25.1, 32.1(a), 40.2(d), 89), data portability rights (Article 20), security by design (Article 32), and more generally technical measures (e.g. Article 17.2). At the same time, the GDPR requires that the controller keeps a proper administration to *demonstrate* compliance (Article 30), departing from the old regime (under the DPD) where controllers had to register their operations with the data protection supervisor.

Next to these novel obligations, the regulation takes enforcement seriously. One of the biggest failures of previous regimes of data protection law was a paramount lack of enforcement, providing no incentive whatsoever to comply. The enforcement chapter of the GDPR, however, provides for a close-knit network of enforcement activities, by individual persons, non-profit organizations, and by the supervisors.

Chapter VIII provides the following enforcement mechanisms, under the heading of ‘Remedies, liability and penalties’:

Articles 77 and 78 provide the data subject with the right to lodge a complaint with a supervisory authority, and the right to an effective judicial remedy against legally binding decisions of a supervisory authority concerning them, including a remedy against the supervisory authority that does not handle their complaint.

Article 79 provides the data subject with direct access to court, apart from their right to lodge a complaint with the supervisory authority. This will enable direct action against the controller or processor, for instance an action for injunctive relief, requesting a court order to prohibit unlawful processing.

Article 80 stipulates that data subjects can mandate their rights from Articles 77–79 to a not-for-profit body, organization or association, enabling such a body to exercise these rights on their behalf. If MS law allows, they can also mandate their right to sue for compensation (based on Article 80). MS law may also provide that a not-for-profit can lodge a complaint with the supervisory authority (as in Article 77) or with the court (as in Articles 78 and 79).

Article 81 regulates suspension of proceedings and jurisdictional issues in the case of simultaneous or overlapping proceedings in different MSs.

Article 82 provides a right to compensation in the case that any person (not just the data subject) has suffered material or non-material damages due to infringements of the GDPR. It stipulates that the controller is liable, adding liability for the

processor if damage is caused by non-compliance with obligations directed specifically to the processor (or by its acting outside or contrary to lawful instructions by the controller). A controller or processor will be exempted from liability if they prove that they are not in any way responsible for the event that caused damage. If more than one controller and/or processor is liable for damage caused, several liability applies (each must pay the full damage, but each can claim back from the others any damage paid beyond their own responsibility). The competent courts will be the same as those competent under MS law for claims based on Article 79.

Article 83 stipulates that supervisory authorities shall impose ‘effective, proportionate and dissuasive’ administrative fines and details the general and specific conditions for such fines. Maximum fines can be €20,000,000, or in the case of an undertaking, up to 4 per cent of the total worldwide annual turnover of the preceding financial year, whichever is higher.

Finally, Article 84 requires that MSs lay down rules for other penalties, in particular for infringements not subject for administrative fines of Article 83.

5.6 Privacy and Data Protection Revisited

In this chapter, we have explored human rights law and investigated the ‘workings’ of the right to privacy in the context of the ECHR, and the right to data protection in the context of the CFREU, as further protected by the GDPR. This cannot be more than a first impression of relevant applicable law. Many relevant provisions and other legislation have not been discussed. The PDPD has not been discussed, the ePD (and its draft successor) have not been detailed. Convention 108 of the Council of Europe has been ignored,²⁷ and a further exploration of the differences between EU and US law has been similarly left aside.

What I hope the reader will take home from this chapter is the salient complexity of privacy and data protection law and the importance of ‘practical and effective’ legal remedies when rights are violated. Though complexity and practical effectiveness may sometimes be incompatible, more often they demonstrate the adaptive nature of legal protection in the face of an increasingly data-driven environment.

²⁷ The 1981 Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data (CETS No. 108), <https://www.coe.int/en/web/data-protection/convention108-and-protocol>.

In Chapter 10 we will return to the subject of EU data protection law with an eye to the increasingly code-driven nature of our environment, highlighting the unique nature of EU data protection rights with regard to automated decisions based on the processing of personal data.

References

On the right to privacy

- Council of Europe, Guide on Article 8 of the European Convention on Human Rights. Right to respect for private and family life, home and correspondence, updated on 31 August 2018. https://www.echr.coe.int/Documents/Guide_Art_8_ENG.pdf.
- Korff, Douwe, 2008. *The Standard Approach under Articles 8–11 ECHR and Art. 2 ECHR*. https://www.pravo.unizg.hr/_download/repository/KORFF_-_STANDARD_APPROACH_ARTS_8-11_ART2.pdf.
- Mowbray, Alastair. 2005. ‘The Creativity of the European Court of Human Rights.’ *Human Rights Law Review* 5 (1): 57–79. <https://doi.org/10.1093/hrlrev/ngi003>.

On the concept of family resemblance

- Biletzki, Anat and Anat Matar. ‘Ludwig Wittgenstein.’ *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/sum2018/entries/wittgenstein/> (the quote in this chapter is taken from this entry).

On freedom from and freedom to

- Berlin, Isaiah. 1969. ‘Two Concepts of Liberty.’ In *Four Essays on Liberty*, edited by Isaiah Berlin, 118–73. Oxford and New York: Oxford University Press.

On data protection law

- European Union Agency for Fundamental Rights (FRA). 2018. ‘Handbook on European Data Protection Law—2018 Edition.’ <http://fra.europa.eu/en/publication/2018/handbook-european-data-protection-law>.
- Journal. *International Data Privacy Law*. <https://global.oup.com/academic/product/international-data-privacy-law-20444001>.
- Kuner, Christopher. 2007. *European Data Protection Law: Corporate Regulation and Compliance*. 2nd ed. New York: Oxford University Press (see updates per chapter at: <http://global.oup.com/booksites/content/9780199283859/updates/>).

On privacy as control over personal information:

Westin, Alan. 1967. *Privacy and Freedom*. New York: Atheneum. (the quotation is taken from p. 7).

On privacy as freedom from unreasonable constraints on identity construction

Agre, Philip E., and Marc Rotenberg. 2001. *Technology and Privacy: The New Landscape*. Cambridge, MA: MIT (quotation taken from p. 7).

On the difference between privacy and data protection

Kokott, Juliane, and Christoph Sobotta. 2013. 'The Distinction between Privacy and Data Protection in the Jurisprudence of the CJEU and the ECtHR'. *International Data Privacy Law* 4 (3), 222–28. <http://idpl.oxfordjournals.org/content/3/4/222.full?sid=a0d12330-d8f3-4387-a7dc-58905c9379a2>.

On data protection by design and on legal protection by design

Hildebrandt, Mireille. 2017. 'Saved by Design? The Case of Legal Protection by Design'. *NanoEthics*, August, 1–5. <https://doi.org/10.1007/s11569-017-0299-0>.

Hildebrandt, Mireille, and Laura Tielemans. 2013. 'Data Protection by Design and Technology Neutral Law'. *Computer Law & Security Review* 29 (5): 509–521.

6

Cybercrime

The more we become dependent upon data- and code-driven environments, the more serious the impact of cybercrime. Whereas individual damage or harm may be remedied by way of private law compensation, substantial damage to critical infrastructure, societal trust, and economic welfare requires a complementary approach that re-establishes and confirms the normative foundations of societal intercourse. To some extent this is the task of administrative law, imposing sanctions for violating legal norms that aim to protect what political theory and legal philosophy call ‘public goods’. In economics, the term ‘public goods’ refers to goods that are non-exclusionary because they cannot be monopolized (such as the air we breathe) and non-rivalrous because usage by the one does not imply less use by another (such as information). In political theory and legal philosophy, the term refers to something that benefits society in general, whether that something is exclusionary, non-rivalrous or neither. We can think of welfare, public health, freedom of expression, universal access to electricity, education, or—at a higher level of abstraction—we can think of a fair distribution of income and access to other goods. In this—non-economic—sense, public goods are closely related to shared values, though the term ‘good’ refers to more than an aspiration or mental preference, as it denotes the actual availability of the good. In legal philosophy, human rights are considered as public goods. The GDPR is an example of an *administrative law* that protects public goods such as privacy, non-discrimination, and freedom of expression. The administrative law approach, however, easily conflates sanctions with paying a fee to exempt oneself from following the law. ‘Speeding on a public road is prohibited, and whoever speeds will be fined’ may turn into ‘speeding on a public road is allowed if one is willing to pay the fine’.

Criminal law is not about paying a price for violating societal norms. In the end it is about censure; it is about holding to account those who seriously flout and diminish respect for their fellow citizens, for the public goods that sustain societal peace, and for the individual flourishing it enables. It is not merely about disrespect for a particular person or those close to them, but about the indirect disintegration of societal trust such disrespect brings about.

Criminal law is about punishing those who negate or ignore the shared normativity that societies thrive on. It is far more than a utilitarian calculus meant to deter a *homo economicus* (the calculating human agent) from violating the law, by imposing costs that hopefully overrule the benefits. Neither is criminal law a way to shame vulnerable agents into ‘behaving’ themselves. Criminal law is about *censure*, about *addressing fellow citizens as responsible agents instead of manipulable pawns*.

The monopoly of violence prohibits private punishment or taking the law in one’s own hands. The internal and external sovereignty of states implies that a government that does not protect its citizens against crime or against other states will lose its footing. Criminal law therefore does not merely provide competences, it also constitutes a task. A government that systematically forsakes punitive interventions when criminal offences are committed, may raise fear about further breaches of the societal contract. This places a heavy burden on governments, as they need to provide safety and trust, without, however, themselves violating safety and trust in the process of defending it. This goes for all criminal law interventions, whether investigatory or punitive.

In the case of cybercrime, competent authorities seem to face a moving target. Technological developments, both on the side of perpetrators and on the side of policing and forensics, often outwit prevalent and tested strategies when trying to deal with the special character of cybercrime. This chapter will first raise the question of what makes cybercrime ‘cyber’, followed by an introduction to the international and supranational legal frameworks that are meant to cope with cybercrime. Finally, we will provide a more detailed analysis of the Cybercrime Convention, including a reflection on the image of the weighing scale where it comes to balancing safety and security against rights and freedoms.

6.1 The Problem of Cybercrime

In the Internet Security Threat Report of 2018, Symantec reports:¹

From the sudden spread of WannaCry and Petya/NotPetya, to the swift growth in coinminers, 2017 provided us with another reminder that digital security threats

¹ Symantec, Internet Security Threat Report 2018, 5–6, available at: http://images.mktgassets.symantec.com/Web/Symantec/%7B3a70beb8-c55d-4516-98ed-1d0818a42661%7D_ISTR23_Main-FINAL-APR10.pdf?

can come from new and unexpected sources. With each passing year, not only has the sheer volume of threats increased, but the threat landscape has become more diverse, with attackers working harder to discover new avenues of attack and cover their tracks while doing so.

According to Symantec, one in thirteen web requests leads to malware; 24,000 is the average number of malicious mobile apps blocked each day; 5.4B WannaCry attacks have been blocked. Compared to 2017, Symantec reports an 80 per cent increase in malware attacks on Macs, a 46 per cent increase in new ransomware variants, a 600 per cent increase in attacks against internet of things (iot) devices, a 13 per cent overall increase in reported vulnerabilities, a 29 per cent increase in industrial control system (ICS) related vulnerabilities, and, finally, an 8,500 per cent increase in coinminer detections.

Though these numbers raise many questions—e.g. as to the distribution of high-impact effects compared to mere nuisance—cybercrime is a major threat for consumers, businesses, law enforcement, national security, and critical infrastructure.

Symantec is focused on *cybersecurity*, which is not the same as *cybercrime*. Cybersecurity is usually defined in terms of confidentiality, integrity, and availability (CIA) of either data, computing systems, or both. We will now first examine what is meant by computer crime and cybercrime and how it relates to cybersecurity.

6.1.1 Computer crime

In the time when computing systems were stand-alone devices, what we now call cybercrime was framed as ‘computer crime’.

In an effort to justify this as a special subdomain of law, it was structured as consisting of:

1. crimes committed *with a computer* (the computer as the *instrument* of crime);
2. crimes committed *against a computer* (the computer as the *target* of the crime);
3. crimes committed *in the context of computers* (‘traditional’ crimes committed in an *environment* where computers play an important role).

The computer ‘as an instrument’ concerned offences such as spam or phishing; the computer ‘as a target’ concerned the use of malware or distributed denial

of service (DDOS) attacks; traditional crimes ‘in the context of computers’ would be digital identity fraud, online copyright infringements, or online child pornography.

Another analytical distinction differentiates between:

1. computer-assisted *traditional crimes*, where the nature of the crime is transformed due to the different nature of online environments; and
2. *new types of crimes*, involving the confidentiality, integrity, or availability of digital data or computing systems (here computer crime overlaps with digital security), involving both crimes with and crimes against a computer.

Whether old or new, the question remains what is ‘the difference that makes a difference’ between existing criminal offences and more recently added computer or cyber offences.

6.1.2 Cybercrime

The rise of the internet and the world wide web, the interconnection between computing systems (the resilient routing of packages across a network of nodes), and the hyperlinking of information across the network (resulting in an unprecedented explosion of content, communication, and metadata), signified the shift from computer crime to cybercrime. We can safely say that we now live in a different world than two decades ago. This is related to the affordances of an unprecedented rise of computing power on the one hand (with the implied miniaturization of the carriers of digital data), and hyperconnectivity on the other (with the ensuing network effects).

This makes cybercrime different across six dimensions of human intercourse, in ways that are highly relevant for the criminal law: distance, scale, speed, distribution, invisibility, and visibility.

1. *distance* is implied in the ability to exercise all kinds of ‘remote control’, ignoring traditional, territorial borders, and thereby, for example, causing major issues for the force of law across different jurisdictions;
2. *scale* is implied in the ability to automate scripts that can affect an enormous amount of other automated systems, that can in turn easily

multiply the reach of a message or malware, thus, for example, enabling massive spam and attacks;

3. *speed* is implied in the combination of an exponential increase in computing power and hyperconnectivity, which, for example, enables the immediate or timed destruction of evidence (even at a distance) and an easy way out of criminal accountability;
4. *distribution* is implied in the networked nature of both the various stacks of the internet, the web, and various application layers, across remote servers (cloud computing) and amalgamated in hardware that combines operating systems, firmware, different software, and applications that have been developed by different teams and companies, while default settings may be changed by the seller, by the buyer (e.g. a service provider), and/or by the end-user, presenting all those involved with seemingly unsurmountable problems in the attribution of responsibility when things go wrong, for example, in the case of self-driving cars or data-driven energy grids;
5. *invisibility* is implied in the differentiation between backend systems that call the shots and frontend systems where end-users are presented with a choice architecture that hides the choices made in the backend, presenting huge issues for the foreseeability of one's actions, for forensics, and for the attribution of causality in the case of harm or other types of damage, as, for example, in the case of manipulative micro-targeting of individual political opinion;
6. *visibility* is implied where the collection, linkability, and inferencing of 'big data' and the wonders of machine learning enable the 'legibility' of end-users in ways that render them vulnerable to, for example, identity theft, invisible nudging or manipulation, blackmailing, extortion, and—in the case of children—grooming.

We can continue the listing and move into corporate espionage, cyberwar, and concerted attacks on critical infrastructure, for example, that of energy supply or democratic institutions. Clearly, states, with their 'traditional' monopoly of violence and their 'traditional' *ius puniendi* (the right to impose public punishment), have been struggling to redefine the borderless and initially lawless realm of 'cyberspace', to combat cybercrime by way of policing, forensics, and judicial cooperation. Because cybercrime does not stop at national borders, states are collaborating at the international and supranational level to come to terms with the transnational nature of cybercrime.

6.2 Cybercrime and Public Law

As discussed above, public law consists of constitutional law, international public law, and administrative law. Constitutional law is relevant for cybercrime to the extent that it determines the right to a fair trial, the criminal law legality principle, and the right to privacy (that is often at stake when states create and apply investigatory competences to combat cybercrime). International public law is relevant for cybercrime because the need to act across territorial borders has resulted in concerted efforts to conclude international treaties on cybercrime. Administrative law is relevant for cybercrime to the extent that supranational legislation on cybersecurity (notably EU directives), imposes duties on Member States (MSs) to align their approach across national borders.

The most important treaty that has been initiated to combat cybercrime across territorial borders is the Cybercrime Convention (CC),² initiated by the Council of Europe (CoE). Within the context of the EU, two directives are relevant, notably the Directive on Attacks against Information System,³ and the Directive on Network and Information Security (NIS) (EU) 2016/1148.⁴

6.2.1 The Cybercrime Convention

The CC was initiated by the CoE, though from the beginning some states outside the CoE were involved, notably the United States, Canada, Japan, and South Africa. To this day, it is the most global treaty on cybercrime thus far concluded. The treaty was signed on 23 November 2001 and entered into force on 1 July 2004, after five states had ratified, including at least three MSs of the CoE (in line with Article 36 CC). In the Netherlands, the treaty entered into force on 1 March 2007, in Japan on 1 November 2012, in the United States on 1 January 2007 (treaties are in force in a contracting state once the treaty itself is in force and after the relevant state has ratified, see section 4.2.1 above). The status of accession and ratification on

² Convention on Cybercrime 2001, Treaty ETS No.185 of the Council of Europe.

³ Directive on Attacks against Information Systems 2013/40/EU, available at: <http://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32013L0040&from=EN>.

⁴ Directive (EU) 2016/1148 on Network and Information Security (NIS), available at: <http://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016L1148&from=EN>.

10 January 2020 was: three signatures that have not yet been ratified and sixty-four ratifications.

The idea of the CC is (1) to agree on new *competences* to investigate cyber-crime, adapted to the intricacies of cyber- as opposed to traditional crimes, and (2) on *joint definitions* of criminal behaviour in cyberspace to make sure that offenders cannot avoid charges by escaping to more lenient jurisdictions, while thus (3) ensuring that *legal certainty* is safeguarded across territorial borders, both with regard to investigatory competences and with regard to what qualifies as criminal conduct, while always (4) preserving a proper level of *legal protection* regarding related human rights and freedoms.

The fact that the CC is international and not supranational law means that whether it has direct effect in MSs depends on whether a MS has a monist or dualist system of recognizing the applicability of international law. In the Netherlands, as discussed above, Article 93 of the Netherlands Constitution makes this dependent on the way international law is formulated (see section 4.2.2 above). Direct effect is only at stake when the content of a treaty addresses citizens by way of granting them rights. The CC, however, does not address citizens. It addresses the MSs, requiring them to implement the content of the CC. This means that the CC lacks direct effect and must first be implemented in national law.

The content of the CC can be summed up briefly as:

Substantive criminal law

- Article 1: definitions
- Articles 2–6: CIA crimes
- Articles 7–8: ‘traditional’ crimes
- Article 9: content crime: child porn
- NB see also First Additional Protocol on racism, ETS 189
- Article 10: copyright violations
- Articles 11–13: ancillary provisions

Procedural criminal law

- Articles 14–15: scope
- Articles 16–21: investigation powers
- Article 22: jurisdiction

International cooperation

- Articles 23–35: extradition, mutual assistance between justice authorities, provisional measures, investigative powers

Other provisions

- Articles 36–38: signature, entry into force

This clearly shows the structure of the CC, highlighting the goal of achieving a similar level of protection against cybercrime on the substantive and the procedural level across national borders.

The fact that the CC lacks direct effect raises the following questions:

1. Can Dutch police base their investigations into cybercrime on the CC?
2. Can a victim of online credit card fraud sue the perpetrator based on the CC?
3. Can a Dutch court convict on the basis of the CC?

The answer should be clear by now: the police cannot base their investigations on legal powers attributed by the CC (only on competences attributed by national law that implements the CC); a victim of online credit card fraud cannot sue the perpetrator based on the CC (the CC does not concern private law, the police and/or the public prosecutor hold the monopoly to initiate a criminal charge); a court cannot convict a perpetrator based on the CC (only based on a criminalization enacted in national law that implements the CC).

6.2.1.1 Substantive law

Note that the CC assumes that the *criminal law legality principle* is in force (see above section 3.1.3 and 3.3.2.1): no punishment without prior and precise criminalization, as, for example, defined in Article 7 ECHR that is binding for the MSs of the CoE. By imposing legal obligations on contracting parties to criminalize specified conduct under the heading of cybercrime, the CC reasserts that criminalization in the legal sense is a prerequisite of fighting cybercrime in constitutional democracies.

The first set of criminal offences concerns CIA-related crimes, such as hacking, or computer trespass, interception, data interference, and system interference. I will discuss them more extensively, as they are highly relevant for computer scientists.

Hacking or computer trespass must be criminalized as stipulated by Article 2:

Article 2—Illegal access

Each Party shall adopt such legislative and other measures as may be necessary to establish as criminal offences under its domestic law, when committed intentionally, the access to the whole or any part of a computer system without right. A Party may require that the offence be committed by infringing security measures, with the intent of obtaining computer data or other dishonest intent, or in relation to a computer system that is connected to another computer system.

The legal effect of this provision is that contracting parties are obligated to enact the relevant criminalization, under the legal conditions specified.⁵ This means that, structured in terms of legal effect and legal conditions, parties should enact that:

The *legal effect* of ‘access to the whole or any part of a computer system’ being a criminal offence, depends on the *following legal conditions*:

- it has been committed intentionally, and
- without right.

Additionally, a party *may require* that to qualify as a criminal offence, such access is achieved:

- by infringing security measures, and/or
- with the intent of obtaining computer data or other dishonest intent, and/or
- in relation to a computer system that is connected to another computer system.

What does this mean for ‘ethical hacking’ (penetration testing to detect security problems)? From an ethical perspective, one can distinguish between a black hat hacker (malicious intent), a white hat hacker (good intent and permission), and a grey hat hacker (good intent but no permission). However, from a legal perspective, if a system is hacked intentionally without permission of the user or owner, the act qualifies as a criminal offence, irrespective of good or bad intent, unless there is another ‘right’ to access, such as a legal competence (e.g. for the police, provided the relevant conditions for the exercise of that competence apply). One could think of

⁵ In the Netherlands this has been implemented in Article 138ab of the Netherlands Criminal Code (NCC).

three ways to prevent punishment for ethical hacking (that is, for grey hat hacking).

First, the public prosecutor may decide *not to prosecute* if they find there is no general interest in prosecuting,⁶ for instance because the hacker followed guidelines of responsible disclosure. Note that penetration testing will fall within the scope of this criminal offence unless one has permission or an assignment to conduct such testing. In some countries, the public prosecutor has no discretion to abstain from prosecution, due to a strict interpretation of the procedural criminal law legality principle (see above section 3.3.2.2).

This brings us to the second way that punishment can be prevented, which would entail that a *legal justification* can be invoked, despite the fact that the hacker had no right to access the system.⁷ Such a defence concerns the requirement that a criminal offence implies ‘wrongfulness’ as part of the *mens rea* that constitutes a criminal offence (see above section 3.3.2.1). One could, for instance, claim that a higher—legally relevant—duty overruled the duty to refrain from intentionally accessing the system without right. It will be up to the court to decide whether such a higher duty justified unlawful access. Note that once a court acknowledges such a higher duty, this would justify all similar cases of unlawful access, unless the decision is overruled by a higher court. As the reader may guess, courts will be cautious in accepting such grounds of justification, due to the consequences of such acceptance.

The third way to prevent punishment could be *conviction without punishment*,⁸ which would be a clear sign that the court does not accept the lawfulness of the access, but nevertheless finds good reason in the circumstances of the offence that was committed to abstain from punishment. Note that in jurisdictions that impose minimum sentences for such an offence, without enabling courts to convict without punishment, this is not an option.

After Article 2 on unlawful access, we have another CIA-related offence in Article 3 on *interception*:

Article 3—Illegal interception

Each Party shall adopt such legislative and other measures as may be necessary to establish as criminal offences under its domestic law, when committed intentionally, the interception without right, made by technical means, of non-public

⁶ In the Netherlands this could be based on Article 349 in relation to Article 167(2) of the Netherlands Code of Criminal Procedure (NCCP).

⁷ In the Netherlands this could be based on Article 352 NCCP and Articles 40, 41(1), 42, 43(1) NCC.

⁸ In the Netherlands this can be based on Article 9a NCC.

transmissions of computer data to, from or within a computer system, including electromagnetic emissions from a computer system carrying such computer data. A Party may require that the offence be committed with dishonest intent, or in relation to a computer system that is connected to another computer system.

In terms of legal effect and legal conditions, this provision requires the following:⁹

The *legal effect* of ‘Interception of non-public transmissions of computer data to, from or within a computer system, including electromagnetic emissions from a computer system carrying such computer data’ *being a criminal offence*, depends on the following *legal conditions*:

- it has been committed intentionally, and
- without right, and
- made by technical means.

A party may require that to qualify as an offence such interception be committed:

- with dishonest intent, and/or
- in relation to a computer system that is connected to another computer system.

Article 4 CC stipulates the criminalization of another CIA-related offence, notably that of data interference:¹⁰

- 1 Each Party shall adopt such legislative and other measures as may be necessary to establish as criminal offences under its domestic law, when committed intentionally, the damaging, deletion, deterioration, alteration or suppression of computer data without right.
- 2 A Party may reserve the right to require that the conduct described in paragraph 1 result in serious harm.

In terms of legal effect and legal conditions, this implies that parties enact that:

The *legal effect* of ‘the damaging, deletion, deterioration, alteration, or suppression of computer data’ *being a criminal offence*, depends on the following *legal conditions*:

⁹ In the Netherlands this has been implemented in Article 139c NCC.

¹⁰ Implemented in the Netherlands in Article 350a NCC.

- it has been committed intentionally, and
- without right.

A party may reserve the right to require that such data interference, to qualify as a criminal offence

- results in serious harm.

Article 5 then stipulates the criminalization of the final CIA-related offence, that of *system interference*:¹¹

Each Party shall adopt such legislative and other measures as may be necessary to establish as criminal offences under its domestic law, when committed intentionally, the serious hindering without right of the functioning of a computer system by inputting, transmitting, damaging, deleting, deteriorating, altering or suppressing computer data.

In terms of legal effect and legal conditions this entails that parties must legislate.

The *legal effect* of ‘the serious hindering of the functioning of a computer system’ being a criminal offence, depends on the following *legal conditions*:

- when committed intentionally,
- without right
- by inputting, transmitting, damaging, deleting, deteriorating, altering, or suppressing computer data.

As indicated above, the CIA-related offences are followed by ‘traditional’ crimes such as identity fraud in Articles 7–8, by content crime, notably child porn in Article 9, and by copyright violations in Article 10. These can be analysed similarly to Articles 2–5.

6.2.1.2 Procedural law

The second part of the CC concerns procedural law, effectively stipulating that specified investigatory powers are attributed to the police and justice authorities: expedited preservation of computer data (traffic and content),

¹¹ Implemented in the Netherlands in Article 138b NCC, and in Article 161 *sexies* NCC if such interference hinders or obstructs data ‘in the general interest’ or causes a ‘disruption in a public telecommunications network or in the execution of a public telecommunication service’.

production orders, search and seizure, and interception (metadata and content data). I will provide an analysis of the production order and the legal power to conduct search and seizure and leave it to the reader to study the legal conditions for lawful interception.

As explained in section 2.2.1, legal norms can be distinguished as either primary or secondary rules. Primary rules regulate human intercourse by way of prohibitions and obligations. Secondary rules constitute competences to legislate, govern, or adjudicate, more generally, they constitute the competence to act. Substantive criminal law, discussed in the previous section, can be understood as a set of secondary rules that impose punitive sanctions when specified primary norms have been violated. The first part of the CC stipulates which primary norms must be protected by way of criminalization. The second part of the CC, regarding procedural criminal law, can be understood as a set of secondary rules that defines under what conditions ‘competent authorities’ are allowed to exercise a set of legal powers that should enable them to combat cybercrime. The second part of the CC thus stipulates what secondary norms must be instituted by the contracting parties in the realm of cybercrime investigation.

Article 18 requires contracting parties to enact legal powers for its competent authorities *to request computer data and subscriber information*.

1. Each Party shall adopt such legislative and other measures as may be necessary to empower its competent authorities to order:
 - a) a person in its territory to submit specified computer data in that person’s possession or control, which is stored in a computer system or a computer-data storage medium; and
 - b) a service provider offering its services in the territory of the Party to submit subscriber information relating to such services in that service provider’s possession or control.
2. The powers and procedures referred to in this Article shall be subject to Articles 14 and 15.
3. For the purpose of this Article the term ‘subscriber information’ means any information contained in the form of computer data or any other form that is held by a service provider, relating to subscribers of its services, other than traffic or content data and by which can be established:
 - a) the type of communication service used, the technical provisions taken thereto and the period of service;
 - b) the subscriber’s identity, postal or geographic address, telephone and other access number, billing and payment information, available on the basis of the service agreement or arrangement;

- c) any other information on the site of the installation of communication equipment, available on the basis of the service agreement or arrangement.

Here we see the legality principle at work (see sections 3.3.1.1 and 3.3.2.2 above), as this stipulates that government authorities can only act if there is a legal basis, while in the case of invasive actions such as criminal investigations these actions require a more detailed legal basis. In other words, whenever government authorities act, they must be ‘competent authorities’, meaning they have been attributed the legal powers for their actions. As discussed in section 3.3.1.1, legal competences have a double function: they both constitute and limit the power they attribute. Article 18 requires parties to attribute specified legal powers, based on the assumption that competent authorities can only act within the confines of the specification that constitutes the power. The second paragraph further asserts this, by referring to Articles 14 and 15, which limit the scope of the investigatory competences and stipulate that relevant safeguards must be in place. We will return to this in section 6.2.2 below.

Let me explain the relevance of a production order in the realm of cybercrime by extensively quoting case law of the European Court of Human Rights (ECtHR), which speaks for itself, nicely demonstrating that even without recourse to the CC, the ECHR implies positive obligations for the contracting parties of the CoE to enact legal competences for the police to give a production order. The case is that of *K.U. v. Finland*.¹² To enable easy reading, I have used some bullet points, without, however, changing the text:

7. On 15 March 1999 an unidentified person or persons placed an advertisement on an Internet dating site
- in the name of the applicant,
 - who was 12 years old at the time,
 - without his knowledge.

The advertisement mentioned his age and year of birth,

- gave a detailed description of his physical characteristics,
- a link to the web page he had at the time,
- which showed his picture, as well as his telephone number, which was accurate save for one digit.

¹² ECtHR, 2 December 2008, Application no. 2872/02 (*K.U. v. Finland*).

In the advertisement, it was claimed that he was looking for an intimate relationship with a boy of his age or older

- ‘to show him the way’.

9. The applicant’s father requested the police

- to identify the person who had placed the advertisement in order to bring charges against that person.

The service provider, however,

- refused to divulge the identity of the holder of the so-called dynamic Internet Protocol (IP) address in question,
- regarding itself bound by the confidentiality of telecommunications as defined by law.

10. The police then asked the Helsinki District Court (käräjäoikeus, tingsrätten)

- to oblige the service provider to divulge the said information pursuant to section 28 of the Criminal Investigations Act (esitutkintalaki, förundersökningslagen; Act no. 449/1987, as amended by Act no. 692/1997).

11. In a decision issued on 19 January 2001, the District Court refused

- since there was no explicit legal provision authorising it to order the service provider to disclose telecommunications identification data in breach of professional secrecy.

The court noted that by virtue of Chapter 5a, section 3, of the Coercive Measures Act (...) and section 18 of the Protection of Privacy and Data Security in Telecommunications Act (...)

- the police had the right to obtain telecommunications identification data in cases concerning certain offences, notwithstanding the obligation to observe secrecy.

However, malicious misrepresentation was not such an offence.

35. The applicant complained under Article 8 of the Convention that

- an invasion of his private life had taken place and that
- no effective remedy existed to reveal the identity of the person who had put a defamatory advertisement on the Internet in his name, contrary to Article 13 of the Convention.

Article 8 provides:

1. Everyone has the right to respect for his private and family life, his home and his correspondence.
2. There shall be no interference by a public authority with the exercise of this right except such as is in accordance with the law and is necessary in a democratic society in the interests of national security, public safety or the economic well-being of the country, for the prevention of disorder or crime, for

the protection of health or morals, or for the protection of the rights and freedoms of others.’

Article 13 provides:

‘Everyone whose rights and freedoms as set forth in [the] Convention are violated shall have an effective remedy before a national authority notwithstanding that the violation has been committed by persons acting in an official capacity.’

41. There is no dispute as to the applicability of Article 8:

- the facts underlying the application concern a matter of ‘private life’;
- a concept which covers the physical and moral integrity of the person (see *X and Y v. the Netherlands*, cited above, § 22).

Although this case is seen in domestic law terms as one of malicious misrepresentation,

- the Court would prefer to highlight these particular aspects of the notion of private life,
- having regard to the potential threat to the applicant’s physical and mental welfare brought about by the impugned situation and to his vulnerability in view of his young age.

42. The Court reiterates that,

- although the object of Article 8 is essentially to protect the individual against arbitrary interference by the public authorities,
- it does not merely compel the State to abstain from such interference:
- in addition to this primarily negative undertaking,
- there may be positive obligations inherent in an effective respect for private or family life (see *Airey v. Ireland*, 9 October 1979, § 32, Series A no. 32).

43. These obligations may involve

- the adoption of measures designed to secure respect for private life even in the sphere of the relations of individuals between themselves.

There are different ways of ensuring respect for private life

- and the nature of the State’s obligation will depend on the particular aspect of private life that is at issue.

While the choice of the means to secure compliance with Article 8 in the sphere of protection against acts of individuals

- is, in principle, within the State’s margin of appreciation,
- effective deterrence against grave acts, where fundamental values and essential aspects of private life are at stake, requires efficient criminal-law provisions (see *X and Y v. the Netherlands*, cited above, §§ 23–24 and 27; *August v. the United Kingdom (dec.)*, no. 36505/02, 21 January 2003; and *M.C. v. Bulgaria*, no. 39272/98, § 150, ECHR 2003-XII).

49. The Court considers that practical and effective protection of the applicant required that

- effective steps be taken to identify and prosecute the perpetrator, that is, the person who placed the advertisement.

In the instant case, such protection was not afforded.

- An effective investigation could never be launched because of an overriding requirement of confidentiality.
- Although freedom of expression and confidentiality of communications are primary considerations and users of telecommunications and Internet services must have a guarantee that their own privacy and freedom of expression will be respected, such guarantee cannot be absolute and must yield on occasion to other legitimate imperatives, such as the prevention of disorder or crime or the protection of the rights and freedoms of others.
- Without prejudice to the question whether the conduct of the person who placed the offending advertisement on the Internet can attract the protection of Articles 8 and 10, having regard to its reprehensible nature, it is nonetheless the task of the legislator to provide the framework for reconciling the various claims which compete for protection in this context.
- Such framework was not, however, in place at the material time, with the result that Finland's positive obligation with respect to the applicant could not be discharged.
- This deficiency was later addressed. However, the mechanisms introduced by the Exercise of Freedom of Expression in Mass Media Act (see paragraph 21 above) came too late for the applicant.

Let us now move to Article 19 CC, which requires contracting parties to enact a power to conduct a *search and seizure of stored computer data*.¹³

1. Each Party shall adopt such legislative and other measures as may be necessary to empower its competent authorities to search or similarly access:
 - a) a computer system or part of it and computer data stored therein; and
 - b) a computer-data storage medium in which computer data may be stored in its territory.
2. Each Party shall adopt such legislative and other measures as may be necessary to ensure that
 - where its authorities search or similarly access a specific computer system or part of it, pursuant to paragraph 1.a,

¹³ E.g. implemented in the Netherlands Code of Criminal Procedure in Articles 125 (search), 125j (extended search), 125k (an order to provide excess, by way of an encryption key or password; not to suspect), 125l (legal privilege), 125m (notification), 125n (duty to delete data), 125o (competence to block data).

- and have grounds to believe that the data sought is stored in another computer system or part of it in its territory,
 - and such data is lawfully accessible from or available to the initial system,
 - the authorities shall be able to expeditiously extend the search or similar accessing to the other system.
3. Each Party shall adopt such legislative and other measures as may be necessary to empower its competent authorities to seize or similarly secure computer data accessed according to paragraphs 1 or 2. These measures shall include the power to:
 - a) seize or similarly secure a computer system or part of it or a computer-data storage medium;
 - b) make and retain a copy of those computer data;
 - c) maintain the integrity of the relevant stored computer data;
 - d) render inaccessible or remove those computer data in the accessed computer system.
 4. Each Party shall adopt such legislative and other measures as may be necessary to empower its competent authorities
 - to order any person who has knowledge about the functioning of the computer system or
 - measures applied to protect the computer data therein
 - to provide,
 - as is reasonable,
 - the necessary information, to enable the undertaking of the measures referred to in paragraphs 1 and 2.
 5. The powers and procedures referred to in this article shall be subject to Articles 14 and 15.

As in Article 18, this competence is restricted by the safeguards required in Articles 14 and 15 (see paragraph 5), thus asserting the legality principle.

Paragraph 4 includes the competence to request a password to access a computing system, or to decrypt content. The reference to the safeguards of the Rule of Law in paragraph 5 may imply that such a request cannot be directed to a suspect, as this could violate the privilege of *nemo tenetur*, that is, the *privilege against self-incrimination*. The full privilege reads *nemo tenetur se ipsum accusare*, or no one is bound to incriminate themselves. The ECtHR reads this privilege into Article 6 ECHR, even though it is not explicitly articulated.¹⁴ It mainly guards against unwarranted compulsion, but it does not provide an absolute right; depending on the severity of the public interest that is at stake, the effectiveness of procedural safeguards and

¹⁴ ECtHR, 25 February 1993, Application no. 10828/84 (*Funke v. France*), paragraph 44.

how the information obtained is to be used, infringements can be justified. Whether the ECtHR would consider a categorical competence to order a suspect to provide a password as a violation of Article 6 ECHR is not clear, but this will probably depend on whether effective legal safeguards and proportionality requirements are in place.

Note that the CC does not impose an obligation on contracting parties to enact a legal power for the police to remotely hack into computing systems. Though this is not prohibited, such enactment is not an implementation of the CC and there is no obligation to enable remote access, unless via an already accessed system, based on paragraph 2.

6.2.1.3 Extraterritorial jurisdiction to enforce or investigate

Another caveat concerns the limitation of access to the territory of the investigating state, first in paragraph 1.b, which limits search to databases on the territory of the relevant contracting party. Second, a search in a remote system via an already accessed system, based on paragraph 2, is restricted to cases where the competent authorities ‘have grounds to believe that the data sought is stored in another computer system or part of it in its territory’. This restriction is based on a fundamental principle of international law, which prohibits extraterritorial jurisdiction to enforce. As discussed in section 4.1 and 4.4, the combination of internal and external sovereignty, which constitutes both international and national law, implies that states respect one another’s territorial integrity as part of their sovereignty. Conducting criminal law investigations on the territory of another state has therefore been outlawed since the famous *Lotus* case of the Permanent Court of International Justice (PCIJ) in The Hague of 1927.¹⁵ In this case it was decided that such extraterritorial enforcement jurisdiction is only permitted in case of permission granted by the state on whose territory the investigations take place. Such permission can be ad hoc, but it can also be based on mutual legal assistance treaties (MLATs). Article 32 CC confirms the prohibition of extraterritorial enforcement jurisdiction:

Article 32—Trans-border access to stored computer data with consent or where publicly available

A Party may, without the authorisation of another Party:

- a) access publicly available (open source) stored computer data, regardless of where the data is located geographically; or

¹⁵ (*France v. Turkey*) (1927) PCIJ, Ser. A, No. 10.

- b) access or receive, through a computer system in its territory, stored computer data located in another Party, if the Party obtains the lawful and voluntary consent of the person who has the lawful authority to disclose the data to the Party through that computer system.

Article 32 CC has given rise to hefty discussions, as some contracting parties have enacted competences to remotely hack into computing systems that may be on foreign territory. The reader can imagine that this particular issue has major implications for the force of law, the practice of international law, and for the defining features of internal and external sovereignty. The jury is still out on where this will end.

Articles 20 and 21 CC require contracting parties to provide legal powers to enable *interception by the competent authorities, of traffic data* (Article 20) and *content data* (Article 21). Note that without such powers, the police would commit a criminal offence and become punishable. The reader is invited to study both articles in detail, dissecting the cumulative and alternative legal conditions that must be fulfilled for interception to be lawful.

6.2.2 Limitations on investigative powers

As indicated above, the legality principle requires that governments act in a way that is not arbitrary, sufficiently foreseeable, proportional, and embedded in adequate safeguards. This includes respect for human rights and a proactive approach to potential risks for democracy and the Rule of Law. As mentioned in the previous section, Article 15 explicitly requires the contracting parties to implement all relevant provisions of the CC in line with the demands of constitutional democracy.

Article 15—Conditions and safeguards

1. Each Party shall ensure that the establishment, implementation and application of the powers and procedures provided for in this Section are
 - subject to conditions and safeguards provided for under its domestic law,
 - which shall provide for the adequate protection of human rights and liberties,
 - including rights arising pursuant to obligations it has undertaken under the 1950 Council of Europe Convention for the Protection of Human Rights and Fundamental Freedoms,
 - the 1966 United Nations International Covenant on Civil and Political Rights, and
 - other applicable international human rights instruments, and
 - which shall incorporate the principle of proportionality.

2. Such conditions and safeguards shall, as appropriate in view of the nature of the procedure or power concerned, inter alia, include
 - judicial or other independent supervision,
 - grounds justifying application, and
 - limitation of the scope and the duration of such power or procedure.
3. To the extent that it is consistent with the public interest, in particular the sound administration of justice, each Party shall consider the impact of the powers and procedures in this section upon the rights, responsibilities and legitimate interests of third parties.

Article 15 basically integrates the case law of the ECtHR (the highest court of the CoE, which initiated the CC) into the CC. Above, in section 5.3.5, we have discussed the legal conditions that must be fulfilled for the justification of infringing measures of secret surveillance, such as notably *Weber & Saravia*.¹⁶ The safeguards stipulated in such case law are highly relevant for the competences that must be attributed by the contracting parties and similarly, and the proportionality test of Article 8 ECHR (see also Article 15, paragraph 1) must be built into the procedures that condition the application of these legal powers.

6.2.2.1 Proportionality test for police access to personal data

An interesting example of a proportionality test regarding police access to personal data retained by internet service providers (ISPs) was conducted by the CJEU in its judgment of October 2018.¹⁷ The case concerned a police request to obtain identifying information on those who interacted with a stolen smartphone during a twelve-day period after the phone was stolen. The question was whether this constituted a ‘serious’ interference with the fundamental rights and freedoms of those persons. The CJEU finds in paragraph 60:

It is therefore apparent that the data concerned by the request for access at issue in the main proceedings only enables the SIM card or cards activated with the stolen mobile telephone to be linked, during a specific period, with the identity of the owners of those SIM cards. Without those data being cross-referenced with the data pertaining to the communications with those SIM cards and the location data, those data do not make it possible to ascertain the date, time, duration and recipients of the communications made with the SIM card or cards in question, nor the locations where those communications took place or the frequency of those

¹⁶ ECHR, 29 June 2006, Application no. 54934/00 (*Weber & Saravia v. Germany*).

¹⁷ CJEU, 2 October 2018, C-207/16 (*Ministerio Fiscal v. Juzgado de Instrucción No. 3 de Tarragona*).

communications with specific people during a given period. Those data do not therefore allow precise conclusions to be drawn concerning the private lives of the persons whose data is concerned.

This was the first step in a proportionality test, weighing the proportionality between the infringement and the purpose it aimed to achieve. Paragraph 61 concludes that the request is not a ‘serious’ infringement. For the proportionality test, the CJEU concludes in paragraph 62:

As stated in paragraphs 53 to 57 of this judgment, the interference that access to such data entails is therefore capable of being justified by the objective, to which the first sentence of Article 15(1) of Directive 2002/58 refers, of preventing, investigating, detecting and prosecuting ‘criminal offences’ generally, without it being necessary that those offences be defined as ‘serious’.

Directive 2002/58 is the ePrivacy Directive which aims to protect the *confidentiality of online communication*. Article 15 of said directive allows for legislative measures that restrict the protection granted in the ePrivacy Directive, for purposes such as prevention and investigation of criminal offences. The CJEU basically states that Article 15(1) does not limit such restrictions to prevention and investigation of ‘serious’ criminal offences. Together with the finding that the request of identifying information that is at stake in this case is not a serious infringement, the CJEU concludes that it is allowed to investigate criminal offences that are not considered serious by way of investigative measures that do not constitute a serious infringement. Though the CJEU is the highest court of the EU and not the highest court of the CoE, its proportionality test is relevant as it follows that of the ECtHR, due to Article 52(3) of the CFREU:

2. In so far as this Charter contains rights which correspond to rights guaranteed by the Convention for the Protection of Human Rights and Fundamental Freedoms, the meaning and scope of those rights shall be the same as those laid down by the said Convention. This provision shall not prevent Union law providing more extensive protection.

6.2.2.2 Proportionality test, balancing tests, and the image of the scale

As a final touch down, I want to briefly discuss the image of the scale that is so often invoked when proportionality comes into focus. In much literature we encounter the idea that security and liberty are mutually exclusive, suggesting

that we can't eat our cake and have it too. This suggests that a *trade-off* between security measures and liberty rights are a given: more of the one supposedly results in less of the other. This is not correct as far as digital security is concerned. Security measures such as encryption will often enable or reinforce a user's capability to make freely chosen and well-informed decisions about sharing personal data. Nevertheless, the opposite is equally incorrect. Some security measures will require disclosure, penetration testing, or even deep packet inspection to facilitate attack monitoring, and this will necessarily infringe individual rights and freedoms, especially where such measures are often invisible or even secret.

The idea that security measures and liberty rights must be framed in terms of a trade-off is not restricted to the domain of cybersecurity. It also pervades the broader domain of policy science where it refers to national and public security, the fight against transnational terrorism and foreign intelligence targeting critical infrastructure and democratic processes. Here, security denotes threats to a person's autonomy and bodily integrity, an organization's resilience, a state's existence or economic welfare, based on targeted attacks. In that sense security is a subdomain of safety, which also refers to threats, though not necessarily based on deliberate targeting.

In the context of cybercrime law, the broader discussion of a trade-off between security and liberties plays out whenever investigatory measures infringe human rights such as privacy, freedom of expression, or, for example, the privilege against self-incrimination. The CC, as we have seen in the preceding subsections, requires proportionality between the infringing measures and the objective that is meant to be protected. It is crucial to acknowledge that such proportionality is not equivalent with the trade-off that is often suggested when the image of the scale is invoked (more protection of security requires less human rights protection), though we should also not take the opposite position that such a trade-off never occurs.

In a seminal article, written shortly after the attacks of 9/11 on the New York World Trade Centre, legal philosopher Jeremy Waldron discussed six caveats for invoking the image of the scale:

1. diminishing liberties does not automatically increase security (a trade-off is not given);
2. 'scale' suggests a precision that is absent, because a *tertium comparationis* is usually absent;

3. liberties cannot be traded at will, they are preconditional for a legitimate government;
4. trading liberty against security often generates a distributive effect (trading the liberty of one group to increase the security of another group);
5. diminishing liberties will increase insecurity in relation to the state;
6. 'scale' has high symbolic value; may contain no effective safeguards whatsoever.

Sometimes, increased security requires infringement of, for example, privacy, but this is not necessarily the case. Some digital security measures may indeed increase privacy protection, for instance when end-to-end encryption is seen as a security measure. In the domain of police investigations into cybercrime, however, whereas the police may use such measures for their internal communication, consumers that employ them may be seen as obstruction of police investigations. Within the context of cybercrime, security measures concern police access to computing systems, production orders, and interception. The first point made by Waldron highlights that the mere fact that such measures infringe privacy does not imply that they increase public security. If a then b does not imply that if b then a.

This also relates to his sixth point: security measures often promise more than they can effectively achieve. In itself this is to be expected, but when a balancing test is done, we must accept that measures that are ineffective cannot be necessary and thus not proportional.

Thinking in terms of a trade-off, using the image of the scale, suggests that the trade-off between liberty and security is a matter of *calculation*: some amount of liberty is traded against some amount of security. Waldron's second point is that this is clearly not the case. Though a security measure may—metaphorically—be understood in terms of costs (liberties) and benefits (security), there is no generally valid way of counting either the costs or the benefits. Security is a different 'thing' than liberty, while both can be understood as public goods as well as private interests. Though one could 'rank' costs and benefits, this does not imply they can be added up or deducted on one and the same scale, which is exactly what the image of the scale lures us into assuming.

This again links to the sixth point; we should not mistake a security measure for the effect it aims to achieve.

The idea of a trade-off also wrongly assumes that liberty and security are independent variables, whereas in a constitutional democracy there are many

dependencies between them. As discussed in section 2.2 and 3.3, in a constitutional democracy a government must not only: ‘(1) act with an eye to the public interest, but also (2) act within the confines of the legality principle and (3) treat citizens with equal respect and concern.’ This entails that liberty is not something to trade at will against other public goods, but something that—like security—is constitutive for a legitimate government. Often, as citizens, we cannot be secure in our life and limbs if liberties can be flouted by the state in its struggle to provide us with security. This connects with the fact that diminishing liberties will often increase insecurity in relation to the state.

Convincing people to give up some liberty to gain some security misrepresents a reality where the liberty of one group of people may be diminished to ensure increased security of another group. The security of some is then traded against the liberty of others. Depending on what kind of security measures are at stake, liberty may be redistributed, for instance, when those dependent on welfare benefits are exposed to automated decision systems that disregard their privacy, whereas others can afford to protect themselves by buying an expensive but well protected smart phone. Data protection law may protect legal residents of the EU whereas illegal aliens may find themselves ‘naked’ in the eye of the immigration machine, seeing their privacy ‘traded’ for the security of already securely settled lawful residents.

6.3 The EU Cybercrime and Cybersecurity Directives

In the strict sense, the Directive on attacks against information systems (2013/40/EU) is ‘the’ EU *Cybercrime Directive*. As to its aims and instrumental value, it overlaps with the CC, requiring EU MSs to criminalize illegal access, attacks against information systems and computer data, and illegal interception (substantive criminal law). Other than the CC it does not concern the criminalization of fraud, child pornography, or copyright violations, clearly focusing on CIA-related offences. Also, other than the CC, it does not impose obligations regarding criminal procedure and criminal investigations. The goal of the directive is *minimum harmonization*, meaning that MSs can go beyond what is required, but not below that, as Article 1 states under the heading of ‘Subject matter’:

This Directive establishes minimum rules concerning the definition of criminal offences and sanctions in the area of attacks against information systems. It also

aims to facilitate the prevention of such offences and to improve cooperation between judicial and other competent authorities.

Interestingly, this directive obligates MSs to impose ‘minimum maximum’ penalties for specific cybercrimes, for example, in Article 9, paragraph 2:

Member States shall take the necessary measures to ensure that the offences referred to in Articles 3 to 7 are punishable by a maximum term of imprisonment of at least two years, at least for cases which are not minor.

Criminal law is often considered core to internal sovereignty, meaning that states resist supranational interference with their criminal law policy. By stipulating *minimum maximum punishment*, EU law reserves discretion for MSs that reject minimum punishment or allow conviction without punishment (as in the Netherlands in Article 9a NCC, see above section 4.1.2 and 6.2.1.1).

The directive pays special attention to criminal law liability for legal persons, and for issues of jurisdiction, and includes various types of cross-national collaboration within the Union (e.g. information exchange via national points of contact, and collection of relevant statistics).

Next to the ‘real’ EU Cybercrime Directive, the EU has also enacted a *cybersecurity directive*, the Directive on Network and Information Security (NIS) (EU) 2016/1148. The subject matter here is defined in Article 1 as:

1. This Directive lays down measures with a view to achieving a high common level of security of network and information systems within the Union so as to improve the functioning of the internal market.
2. To that end, this Directive:
 - (a) lays down obligations for all Member States to adopt a national strategy on the security of network and information systems;
 - (b) creates a Cooperation Group in order to support and facilitate strategic cooperation and the exchange of information among Member States and to develop trust and confidence amongst them;
 - (c) creates a computer security incident response teams network (‘CSIRTs network’) in order to contribute to the development of trust and confidence between Member States and to promote swift and effective operational cooperation;
 - (d) establishes security and notification requirements for operators of essential services and for digital service providers;

- (e) lays down obligations for Member States to designate national competent authorities, single points of contact and CSIRTs with tasks related to the security of network and information systems.
- (...)
6. This Directive is without prejudice to the actions taken by Member States to safeguard their essential State functions, in particular to safeguard national security, including actions protecting information the disclosure of which Member States consider contrary to the essential interests of their security, and to maintain law and order, in particular to allow for the investigation, detection and prosecution of criminal offences

The NIS Directive is not about criminal law, which—as Article 2, paragraph 6 demonstrates—may even be at odds with the information exchange that is at the heart of the NIS Directive. This further clarifies that cybersecurity and cybercrime should not be confused. The CIA-related cybercrimes basically concern the criminalization of attacks on cybersecurity, such as illegal access, attacks on information systems, and illegal interception. In that sense, the NIS Directive overlaps with the CC and the Cybercrime Directive in its objective of identifying, preventing, and deterring threats to cybersecurity.

Note that Article 2 of the NIS Directive states that personal data processed pursuant to this directive falls within the scope of the Data Protection Directive (now the GDPR), meaning it does not fall within the scope of the Police Data Protection Directive (which is focused on processing of personal data ‘by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data’). This, again, clarifies the difference between cybersecurity and cybercrime. The concept of ‘security of network and information systems’ is actually defined in Article 4(2) of the NIS Directive:

‘security of network and information systems’ means the ability of network and information systems to resist, at a given level of confidence, any action that compromises the availability, authenticity, integrity or confidentiality of stored or transmitted or processed data or the related services offered by, or accessible via, those network and information systems;

By defining security in terms of resilience against ‘any action that compromises the availability, authenticity, integrity or confidentiality’, the NIS Directive anchors itself in the CIA concerns that define cybersecurity.

References

On cybercrime law

- Brenner, Susan W. 2012. *Cybercrime and the Law*. Boston: Northeastern University Press.
- Schwarzenegger, Christian, Finlay Young, Gian Ege, and Sarah J. Summers. 2014. *The Emergence of EU Criminal Law: Cyber Crime and the Regulation of the Information Society*. Oxford and Portland: Hart Publishing.
- Tosoni, Luca. 2018. 'Rethinking Privacy in the Council of Europe's Convention on Cybercrime'. *Computer Law & Security Review*, September. <https://doi.org/10.1016/j.clsr.2018.08.004>.

On production orders

- Hert, Paul de, Cihan Parlar, and Juraj Sajfert. 2018. 'The Cybercrime Convention Committee's 2017 Guidance Note on Production Orders: Unilateralist Transborder Access to Electronic Evidence Promoted via Soft Law'. *Computer Law & Security Review* 34 (2): 327–36. <https://doi.org/10.1016/j.clsr.2018.01.003>.

On extraterritorial jurisdiction to enforce in cyberspace

- Hildebrandt, Mireille. 2013. 'Extraterritorial Jurisdiction to Enforce in Cyberspace? Bodin, Schmitt, Grotius in Cyberspace'. *University of Toronto Law Journal* 63 (2): 196–224. <https://doi.org/10.3138/utlj.1119>.

On the image of the scale

- Hildebrandt, M. 2013. 'Balance or Trade-off? Online Security Technologies and Fundamental Rights'. *Philosophy & Technology* 26 (4): 357–79.
- Waldron, Jeremy. 2003. 'Security and Liberty: The Image of Balance'. *Journal of Political Philosophy* 11 (2): 191–210. <https://doi.org/10.1111/1467-9760.00174>.

7

Copyright in Cyberspace

For computer scientists, the most relevant part of copyright law concerns copyright on computer programs, or software. In this chapter, I will provide an introduction to the domain of intellectual property (IP) rights, of which copyright is one—important—example. Before zooming in to copyright on software, which is the enabling precondition for the General Public Licence (GPL) and the open source initiative, I will first investigate the position of IP law in the context of constitutional democracy and clarify that IP law is private law.

In Chapter 2 we discussed the role of law and the rule of law as *both constitutive and limitative*. The constitutive role can be found in the ordering of an architecture that enables individuals to flourish in society, aiming for legal certainty, justice, and the purposiveness of the law. Part of this architecture is the creation of an incentive structure, such as an economic market, that stimulates transactions that in turn stimulate productivity in terms of products and services. The rule of law entails that the legal powers that constitute this architecture are simultaneously restricted, thus introducing the checks and balances of a fair and open economic market.

The attribution of IP rights supposedly creates an incentive structure for the creation of works, technical inventions, trademarks, and designs. Because these rights provide both (1) control and (2) the ability to reap financial rewards for their authors or inventors, they incentivize the creation of these immaterial goods, which would otherwise not be protected. The reason is that if such control were not created by law, it would be very hard to protect one's work or invention other than by keeping it secret.

Other than tangibles and real estate intellectual goods are neither rivalrous nor exclusionary; one person accessing a work does not imply that others can no longer access it, and one person applying an invention does not reduce the extent to which another person can apply it.

The limitative role that aims to protect citizens as deserving equal respect and concern can be traced in the fact that IP rights differ from other property rights because:

1. they are limited in time;
2. after a specified time period the work or invention will enter the public domain.

Basically, the creation of IP rights ensures that the right-holders have an incentive to share their work or invention due to the rewards they will obtain, while simultaneously ensuring that such work or invention will become available for all after a certain period of time. IP law has a limitative role also in that it imposes certain restrictions on the exercise of IP rights, such as the common law doctrine of 'fair use' and the European doctrine of the home copy.

7.1 IP Law as Private Law

In section 3.1.1, we discussed the difference between absolute and relative rights in private law. We did this by framing a legal system as an architecture of legal relationships, both between the state and its legal subjects, and between those legal subjects. Ownership is an absolute right in the sense that it can be enforced against anybody, as everybody has to respect the right of the owner. A contract creates relative rights that can only be enforced against one's contracting party or parties, since contract is based on consent and only binds those who consented to enter into the contract. A contract thus creates two relative rights that the other party complies with the terms of the contract, for example, in the case of a contract of sale: (1) one party has the right that the other party pays the price, while (2) the other party has the right that the good is delivered or the service provided.

IP law provides authors, inventors, and other IP right-holders with a property right in the 'intellectual good' they created or invented (or, for instance, bought). IP law thus is private law, even if international public law plays an important role in requiring states to protect IP rights. The category of property rights includes ownership but also freehold, leasehold, servitude, right of superficies, apartment right, usufruct, pledge and mortgage, and, finally, a defined and limited set of IP rights (see section 3.2.2). The set of property rights is a closed set, a new type of property right cannot be created at will by way of contract. Since a property right can be enforced against anybody, everybody

must be clear as to which property rights can be enforced against them. This requires some form of publicity. In the case of tangibles, this is achieved by way of possession, in the case of real estate and some IP rights, by way of registration. In the case of copyright, publicity is achieved by the duty to credit the author of a work and by the use of the © sign to indicate the copyright holder (which may be another legal subject than the author, for instance an employer or a publisher). Note that the © sign does not constitute copyright, it merely publicly claims the copyright.

Copyright thus provides the freedom to dispose of one's *absolute right* in a work, even though—as discussed in Chapter 3—absolute rights are limited by written and unwritten law. For instance, abuse of the legal power inherent in a property right can constitute a tort action. An example could be the use of a patent in a medical drug in a way that results in physical harm to those who cannot afford it. As mentioned above, copyright is also limited by the 'fair use' or the 'home copy' exception.

A copyright *licence*, however, is a relative right. Whoever has the absolute right in a 'work', can decide to license another to exercise some of their rights, for example, the right to reproduce and publish. This creates a legal relationship between the licensor and the licensee. As discussed in section 3.1.2, in the case of contract, the freedom to contract is default, but nevertheless limited by written and unwritten law. Based on the Software Copyright Directive, for instance, the licensor cannot stop the licensee from creating a backup if that is necessary to run the program. Based on the case law of the Court of Justice of the European Union (CJEU), the licensor cannot stop the licensee from 'performing acts necessary to observe, study or test the functioning of the program, as long as these acts do not infringe the copyright in that program'.¹

Another relative right that may be at stake in the case of copyright is the right to compensation or the right of injunctive relief in the case of a tort. This is highly relevant when third parties violate a right-holder's copyright by illegally uploading or downloading works to which they have no right. Such tort liability would require the right-holder to provide evidence of wrongfulness, damage, and causality (if suing for compensation), whereas the defendant may claim that they had an excuse (e.g. that they were not aware that they were violating the copyright as they wrongly but reasonably assumed it was put in the public domain by or on behalf of the right holder). Even more topical, a right-holder may sue an intermediary for wrongfully facilitating others that

¹ CJEU, 2 May 2012 C-406/10 (*SAS v. World Programming*).

violate the copyright. Think of *The Pirate Bay* and the many attempts to obtain court orders prohibiting *The Pirate Bay* from enabling to upload and download protected content, or court orders that impose an obligation on internet service providers (ISPs) to filter and/or block all traffic to *The Pirate Bay*. Note that copyright-holders are often officially represented by a legal person that has the statutory goal of defending their rights, as they often lack the means to file complaints by themselves.

The interaction between the absolute and relative rights in a ‘work’ can be complex. For instance, as I will explain in section 7.5 below, a GPL obliges those who share a work to only share under the conditions of the specific licence. If a licensee violates this obligation, the licensor could sue for breach of contract, but if there was no copyright, the original licensor would not be able to enforce the relevant conditions against third parties.

Precisely because copyright is an absolute right, it will follow the protected good (*droit de suite*), irrespective of subsequent licensing. In that sense, the *protection* of ‘free software’ that is offered by a GPL, *depends on the property right* that underpins the licence.

7.2 Overview of IP Rights

As related above, there is a closed set of IP rights. Though this chapter focuses on copyright, this section offers a succinct—incomplete—overview of the various types of IP that are most relevant. Note that neither a discovery nor an idea can be protected; protection is limited to the ‘expression of an idea in a certain medium’, or to an ‘invention’, which is not merely a discovery.

7.2.1 Copyright

Generally speaking, two types of copyright must be distinguished. First, the moral right of the author to be credited with their authorship; this is an absolute right in the sense that it can be enforced against anyone and everyone, and on top of that it is not transferable. Second, two types of economic rights, (1) an absolute right that can be enforced against anyone and everyone and is transferable, thus enabling the sale or licensing of a copyright, and (2) a relative right that can only be enforced against the licensor, thus allowing the use of, for example, a software licence.

Sometimes, the difference between the absolute and the relative right becomes blurred. In *Oracle v. UsedSoft*,² the CJEU determined that the sale of a copy of a software program with unlimited usage rights must be qualified as a sale (transfer of ownership) even if the contract speaks of a licence to use.

Copyright is the exclusive right to reproduce, distribute, and publish a ‘work’ created by an author.

Note that in continental European law, the right is called ‘author’s right’ instead of ‘copyright’. Section 7.3 will explain the difference. Copyright is limited in time and jurisdiction; the time period is determined in national jurisdiction. A copyright concerns ‘an expression in a particular medium’. In most jurisdictions, the right is attributed automatically once the ‘work’ is created.

Besides the copyright in a ‘work’ in general, we have specific copyrights in ‘design’, in ‘databases’, and in ‘software’ (note that in the United States one can obtain a patent in software, whereas in Europe that is only possible if the software is part of an ‘invention’ that has a material component).

Sections 7.3 and 7.4 will further explore copyright.

7.2.2 Patents

Like copyright, a patent is limited in time and jurisdiction. The time period during which a patent is valid again depends on national jurisdiction.

To qualify as a patent the intellectual good must be:

1. an invention,
2. that is novel, and
3. has an industrial application.

In other words, to qualify as patentable, these three legal conditions must apply (they are cumulative). The legal effect of being granted a patent is that the right-holder has the exclusive right to commercial exploitation, though only if a fourth legal condition has been fulfilled, which stipulates that an

² CJEU, 3 July 2012, C-128/11 (*Oracle v. UsedSoft*).

application must be filed with disclosure of the novel invention and its industrial application. Whoever is the first to apply obtains the patent. Here, the patent differs from copyright as this is automatically attributed when a work has been created; a patent is not automatically attributed when an invention is both novel and has an industrial application. The patent office that registers the application will not investigate whether the applicant is ‘really’ the person who invented the application and neither will it register two patents if a second person proves that they ‘really’ invented the same application. If one considers investing in the industrial application of an invention it is wise to first check whether a patent has already been registered, as, in that case, the commercial exploitation will be an exclusive right of whoever registered first. If one finds that an invention has been ‘stolen’ and registered, the only way to obtain compensation is to initiate a tort action. Note that oftentimes more than one patent applies to one industrial application. In that case, each right-holder can prohibit commercial exploitation by others, meaning that whoever wishes to develop a business model for the application will have to secure permission from all other right-holders. In that sense, a patent includes a ‘right to exclude’ that overrides another’s ‘right to exploit’ (the same goes for joint copyright).

As the reader will imagine, issues with patent law concern the following:

1. What qualifies as an invention? Is the ‘discovery’ of the genetic profile of a specific gene an invention or a discovery? Or should we rephrase and ask under what conditions a genetic profile qualifies as an invention instead of a discovery? Under US patent law, the following is not considered ‘patentably subject matter’: natural laws, phenomena, or products or abstract ideas. In a famous US case in 1980,³ the Supreme Court found that a recombinant—man-made living—bacteria was patentable, because it was man-made (isolated, purified, and its DNA recombined). In a more recent US case of 2013,⁴ the Supreme Court found that ‘a naturally occurring DNA segment is a product of nature and not patent eligible merely because it has been isolated, but that cDNA is patent eligible because it is not naturally occurring’.
2. When should an invention be considered novel? In US patent law, novelty requires that the invention is new, has utility, and is non-obvious.
3. What counts as an industrial application?

³ Supreme Court USA, 16 June 1980, 447 US 303 (*Diamond v. Chakrabarty*).

⁴ Supreme Court USA, 13 June 2013, 133 S. Ct. 2107 (*Association for Molecular Pathology v. Myriad Genetics, Inc.*).

7.2.3 Trademark

A trademark has been defined as ‘signs fit to identify and distinguish a product or service’ and the IP right to a trademark entails the exclusive use of such signs. As with a patent, deposition of the trademark is required.

7.3 History, Objectives, and Scope of Copyright Protection

As discussed in sections 1.3, before the advent of the printing press, writing was done by hand; written text concerned written manuscripts that were painstakingly copied by hand in monasteries and early universities. In those circumstances, there was a clear difference between an ‘authentic’ original and an ‘authentic’ copy. Whereas paper was first invented in China by Cai Lun in 105, and woodblock printing was developed in China between the seventh and ninth centuries, the invention of the movable type technology that became the prevailing technology of the printing press is usually dated to Gutenberg’s invention in 1450.

The proliferation of identical copies of printed text, as a result of the movable type printing press, gave rise to a curious combination of *censure* and *privileges*. Printed text was used for manifestos and pamphlets, with opinions on politics and religion not necessarily to the liking of the church or the sovereign. The urge to control such proliferation of information drove both the Catholic Church and the sovereigns to a variety of measures meant to restrict and control the dissemination of printed text and the ideas they expressed. The Catholic Church, for example, compiled an Index of printed matter that was considered sinful and therefore prohibited (censure). The sovereigns took to the granting of privileges or monopolies to a limited set printers (publishers), based on their willingness to censure the content in line with royal instructions. The struggle to resist and outlaw such censure on printed text took two to three centuries, finally resulting in a shift from royal privileges for a printers’ guild to a subjective right of the author. The protected object of this right is the ‘work’. This is neither an abstract idea (like a mathematical formula), nor the material carrier on which an idea has been expressed (a copy of a printed book), but only refers to the expression itself.

In the course of the eighteenth century, the idea that the choice to disseminate printed text should be in the hands of the author instead of the printer or publisher, let alone the government, was consolidated in Acts of Parliament

and Royal Decrees that established the right of the author to the works they create: in England, the Queen Anne Act of 1710; in France, the Royal Decrees of 1777; and, in the United States, the Federal Copyright Act of 1790. These Acts and Decrees were the first to attribute authors' rights or copyrights, providing authors with the means to obtain remuneration for their intellectual 'work', while also ensuring that after a fixed period of time, the 'work' would enter 'the public domain'.

In copyright law, 'the public domain' has a very specific meaning, referring to all those 'works' to which no exclusionary rights are applicable, meaning that 'the public' is free to access, reproduce, distribute, and further publish such 'works'.

The rise of modern copyright law thus coincided with human rights such as the freedom of expression, which includes the freedom of information.

Finally, in the nineteenth century, 166 contracting parties (states) negotiated a treaty, the Berne Convention of 1886,⁵ to enable the protection of copyright beyond national borders, based on a unified definition of the scope of protection. Though copyright law is national law, the Berne Convention and a number of other treaties aimed to harmonize IP protection and to ensure a certain degree of transnational protection. Article 2 of the Berne Convention defines the object of protection, or 'protected works' as follows:

- (1) The expression 'literary and artistic works' shall include every production in the literary, scientific and artistic domain, whatever may be the mode or form of its expression, such as books, pamphlets and other writings; lectures, addresses, sermons and other works of the same nature; dramatic or dramatico-musical works; choreographic works and entertainments in dumb show; musical compositions with or without words; cinematographic works to which are assimilated works expressed by a process analogous to cinematography; works of drawing, painting, architecture, sculpture, engraving and lithography; photographic works to which are assimilated works expressed by a process analogous to photography; works of applied art; illustrations, maps, plans, sketches and three-dimensional works relative to geography, topography, architecture or science.
- (2) It shall, however, be a matter for legislation in the countries of the Union to prescribe that works in general or any specified categories of works shall not be protected unless they have been fixed in some material form.

⁵ Berne Convention for the Protection of Literary and Artistic Works 1886.

- (3) Translations, adaptations, arrangements of music and other alterations of a literary or artistic work shall be protected as original works without prejudice to the copyright in the original work.
- (4) It shall be a matter for legislation in the countries of the Union to determine the protection to be granted to official texts of a legislative, administrative and legal nature, and to official translations of such texts.
- (5) Collections of literary or artistic works such as encyclopaedias and anthologies which, by reason of the selection and arrangement of their contents, constitute intellectual creations shall be protected as such, without prejudice to the copyright in each of the works forming part of such collections.
- (6) The works mentioned in this Article shall enjoy protection in all countries of the Union. This protection shall operate for the benefit of the author and his successors in title.
- (7) Subject to the provisions of Article 7(4) of this Convention, it shall be a matter for legislation in the countries of the Union to determine the extent of the application of their laws to works of applied art and industrial designs and models, as well as the conditions under which such works, designs and models shall be protected. Works protected in the country of origin solely as designs and models shall be entitled in another country of the Union only to such special protection as is granted in that country to designs and models; however, if no such special protection is granted in that country, such works shall be protected as artistic works.
- (8) The protection of this Convention shall not apply to news of the day or to miscellaneous facts having the character of mere items of the press.

In the course of the twentieth century, IP protection became an essential part of trade agreements, resulting in an important role for the World Trade Organization (WTO), and in the Trade-Related Aspects of Intellectual Property Rights (TRIPs) Agreement that aims to establish a global regime of protection. In this book, the focus will be on EU law, with some attention to US law, taking note that IP treaties are international public law even if their subject matter is private law and private law is mainly national law. This is also the reason to focus on EU jurisdiction, as it directs national law regarding copyright in the Member States (MSs).

Before investigating EU ‘copyright law’, I will first discuss briefly how UK and US copyright law differs from continental European ‘authors’ law’.

In the continental European tradition, the focus has been on the author and the work. This understanding of ‘authors’ law’ built on the Age of Romanticism of the eighteenth and nineteenth centuries, where the singularity of creative imagination of an individual author took precedence over the mundane business

interests of a publisher. The idea was that ‘authors’ law’ is part of ‘natural law’ rather than being ‘posited’ by a legislator (*positive law*). The ‘authors’ right’, in that line of thinking, is constituted by the original act of creation of the author and should not be tied to formalities (such as registration), while the ‘work’ that is created belongs to the ‘author’s domain’. This is a matter of personality rights (*droit moral* or moral right), rather than a matter of ownership (as Locke would have it).

In the *common law* that inspired the United Kingdom and the United States, the focus was not on the author and their work, but on the original and the copy. This was less a matter of personality and romantic imagination than a matter of pragmatism. Copyright was simply a choice made by a legislator (*positive law*), rather than a natural right inherent in the author’s act of creation. This led to the requirement of registration and an emphasis on copyright as an economic, not a moral right. Here, copyright law is about the domain of the ‘work’ rather than the domain of the ‘author’, and such work is considered original in the sense of not being copied, rather than original in the sense of being creative or novel.

Despite these differences, copyright law—as well as patent law—is generally found to have four objectives, both in the realm of Anglo-American law and in the realm of continental European law:

1. to reward the author or inventor;
2. to provide the author and the inventor with exclusionary control over the use others can make of their work or invention;
3. to incentivize investment in creative expression, invention, and innovation;
4. to ensure the societal benefit of having such expression or invention in the public domain after a fixed period of time.

There is a certain tension between these goals and many IP scholars wonder to what extent current developments undermine the balancing act required to sustain all these goals. For instance, Dusollier writes, on the rhetoric of ‘remuneration-based or control-centered’ models of copyright and patent:

If that rhetoric is revealed as merely one choice amongst others, the imperative of making copyright or patent right an increasingly stronger instrument of control may well be undermined, which could ultimately resignify the meaning of intellectual property.

As to the scope of a copyright, we can sum up the following control-rights:

1. publication (communication to the public);
2. reproduction (making a copy);
3. distribution (of the tangible original or copy);
4. right to prohibit 1, 2, 3;
5. right to license others to exercise the rights of 1, 2, 3.

Note that the right-holder's exclusive right to distribution is exhausted after first sale, thus enabling the sharing of individual copies of books (and the sale of second-hand books). Under the EU Copyright Directive this does not apply to an ebook, as the right to distribution only concerns tangible copies.⁶ This is different in the case of a computer program that falls within the scope of the EU Software Copyright Directive, which should be interpreted as stipulating that the copyright is exhausted after first sale, even if software has been downloaded instead of being supplied on a tangible carrier.⁷

7.4 EU Copyright Law

As emphasized above, copyright law is private law and part of the national jurisdiction of the MSs. There is no European private law, even though various types of international treaties aim to solve cross-jurisdictional problems as part of international private law (see section 4.1.1). There is also no EU private law, even though exceptions apply when the EU legislature imposes duties on MSs to integrate private law liability into their national jurisdiction (e.g. in environmental and data protection law), which the CJEU often interprets as having an autonomous meaning based on the relevant EU directive or regulation.⁸

The EU copyright framework aims to harmonize the applicable law in the MSs, in order to ensure equivalent protection of the right-holders within the confines of the internal market, thus stimulating economic transactions

⁶ In C-263/18 the CJEU decided that making available of ebooks to subscribers must be qualified as 'communication to the public' and does not count as 'distribution' under the Copyright Directive. See section 7.4 for clarification.

⁷ In CJEU, C-128/11 (*Oracle/UsedSoft*), the court determined that the sale of a copy of a software program with unlimited usage rights must be qualified as a sale (transfer of ownership) even if the contract speaks of a licence to use. See section 7.4 for clarification.

⁸ CJEU, 12 March 2002, C-168/00 (*Leitner*); CJEU, 25 October 2005, C-350/03, (*Schulte*); and CJEU, 2 June 2005, C-229/04 (*Crailsheimer Volksbank*).

across national borders within the EU. Recently, the Copyright Directive has undergone a major update,⁹ involving—amongst others—the imposition of two highly controversial legal obligations on ISPs (paraphrased as a ‘link-tax’ and a ‘private policing’ obligation, which I will briefly discuss under section 7.4.1.6).

7.4.1 The Copyright Directive and the Enforcement Directive

The legal framework of EU copyright law is based upon the EU Copyright Directive and the IP Enforcement Directive.¹⁰ Both the EU Software Copyright Directive, and the EU Database Directive are a *lex specialis*,¹¹ meaning that they provide more specific legislation for the copyright in software and that in a database. A *lex specialis* has priority over the *lex generalis* (a law with more general application), just like *lex posterior derogat legi priori*, meaning that a more recent law has precedence over previous legislation.

7.4.1.1 The scope of protection (restrictions) and the limitations

The Copyright Directive requires MSs to offer *the following scope of protection*: the exclusive right to authorize or prohibit reproduction of a work (Article 2); the exclusive right to authorize or prohibit publication or ‘communication to the public’ of a work (Article 3); and the exclusive right to authorize or prohibit the right of distribution of an original or a copy of the work (Article 4), noting that the right of distribution is exhausted after first sale with the consent of the copyright-holder (Article 4.2). Interestingly, the directive also provides legal protection against circumvention of technological measures (think of digital rights management technologies, or DRM) (Articles 6 and 7).

The *scope of protection* is formulated in terms of *restrictions*, as the right allows the copyright-holder to restrict others from reproducing, publishing, or distributing a work without permission.

⁹ Directive 2001/29/EC of the European Parliament and of the Council of 22 May 2001 on the harmonization of certain aspects of copyright and related rights in the information society, and Directive (EU) 2019/790 (on copyright and related rights in the Digital Single Market and amending the copyright and the database directives).

¹⁰ Directive 2001/29/EC (copyright), as amended by Directive (EU) 2019/790, and Directive 2004/48/EC (enforcement IP).

¹¹ Directive 2009/24/EC (copyright software), and Directive 96/9/EC (Database Directive), as amended by Directive (EU) 2019/790.

Limitations are defined in Article 5. These concern, for instance, reproduction, distribution, and possibly publication for teaching, scientific research, caricature, parody, or pastiche. A limitation, here, means that the said right to restrict others is limited.

An important limitation of copyright applies to ‘private use’, defined in Article 5.2(b) as:

Reproduction (...) by a natural person for private use and for ends that are neither directly nor indirectly commercial, on condition that the rightholders receive fair compensation

In Article 5.5 the Copyright Directive stipulates that limitations are only valid if they comply with the so-called ‘triple test’:

Limitations shall only be applied in

- certain special cases
- which do not conflict with a normal exploitation of the work or other subject-matter and
- do not unreasonably prejudice the legitimate interests of the rightholder.

This triple test is similar to the test required by the Berne Convention (Article 9.2), ensuring that limitations of copyright will remain exceptions and do not hinder the normal exploitation, while respecting the economic interests of the right-holders.

Article 5.5 has been relevant in pivotal case law concerning the exception regarding home copies (private use), more precisely with respect to the question of whether downloading of unlawfully uploaded copies of, for example, movies or songs falls within the confines of Article 5.2(b), or must be said to violate Article 5.5.

7.4.1.2 The home copy case of the CJEU

In a landmark case, the Netherlands Supreme Court submitted the following preliminary questions to the CJEU:¹²

Is the exception [mh: of private use, Article 5.2(b)] valid irrespective of whether or not the home copies were reproduced illegally, or is the exception only valid in case the copies were reproduced legally?

¹² CJEU, 10 April 2014, Case C-435/12.

If the last, can the criteria of 5.5 be used to extend the scope of the exception or can it only be used to restrict that scope?

If the last, does fair compensation in the case of home copies of illegally reproduced works violate 5.5 or any other EU legal norm? Is the fact that technical means to effectively enforce a prohibition of illegal downloading relevant for the answer to this question?

The court ruled that:¹³

EU law, in particular Article 5(2)(b) of Directive 2001/29/EC (...), read in conjunction with paragraph 5 of that article, must be interpreted as

- precluding national legislation,
- such as that at issue in the main proceedings,
- which does not distinguish the situation in which the source from which a reproduction for private use is made is lawful from that in which that source is unlawful.

This meant that downloading of unlawfully uploaded content was a violation of copyright that does not fall within the scope of the home copy exception. Thereby it became clear that downloading unlawfully uploaded content was wrongful and potentially constitutes a tort, though not necessarily a criminal offence (as this requires an explicit and prior criminalization of the relevant wrongful conduct).

7.4.1.3 IP enforcement against intermediaries

With regard to the enforcement of copyright, the Copyright Directive requires that in case of infringements, MSs foresee sanctions that are 'effective, proportionate and dissuasive' (Article 8.1), enable 'action for damages and/or apply for an injunction and, where appropriate, for the seizure of infringing material devices, products or components' (Article 8.2), and make sure 'that rightholders are in a position to apply for an injunction against intermediaries whose services are used by a third party to infringe a copyright or related right' (Article 8.3).

The IP Enforcement Directive also applies. Recital (23) of that directive stipulates that right-holders should be entitled to an injunction against intermediaries. This is further elaborated in Article 8 that provides a 'Right of information', obliging MSs to ensure that judicial authorities may order that

¹³ Ibid. para. 66.

information is provided on infringement related actions and actors. On top of that, Article 9 foresees ‘Provisional and precautionary measures’, such as an interlocutory injunction against an intermediary to prevent further infringement.

This is where the eCommerce Directive becomes relevant,¹⁴ first of all, because Articles 12, 13, and 14 provide that providers of ‘mere conduit’, ‘caching’, and ‘hosting’ are not liable for the content shared by third parties on their platform, unless they have been notified that such content violates, for example, copyright. This exemption of liability, however, ‘shall not affect the possibility for a court or administrative authority, in accordance with Member States’ legal systems, of requiring the service provider to terminate or prevent an infringement (art. 12.3, 13.3 and 14.3)’. In the case of a hosting company, the exemption of liability moreover does not ‘affect the possibility for Member States of establishing procedures governing the removal or disabling of access to information (art. 14.3)’.

Next to these exemptions of liability for a subset of ISPs, and the fact that they can nevertheless be ordered to end access or infringement, Article 15 establishes a prohibition to impose a ‘general obligation to monitor’:

1. Member States shall not impose a general obligation on providers, when providing the services covered by Articles 12, 13 and 14, to monitor the information which they transmit or store, nor a general obligation actively to seek facts or circumstances indicating illegal activity.
2. Member States may establish obligations for information society service providers promptly to inform the competent public authorities of alleged illegal activities undertaken or information provided by recipients of their service or obligations to communicate to the competent authorities, at their request, information enabling the identification of recipients of their service with whom they have storage agreements.

The confluence of a duty to foresee that ISPs can be required to remove infringing content with a prohibition to impose a general monitoring obligation has given rise to interesting case law, balancing the right to property and an effective remedy to enforce the property right (of copyright-holders), against the right to conduct a business (of ISPs that may have to pay for filtering and blocking of infringing content) and against the rights of privacy and data

¹⁴ Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular electronic commerce, in the Internal Market (‘Directive on Electronic Commerce’).

protection (of data subjects whose online behaviour may be monitored to detect infringement).

7.4.1.4 Injunctions to cease unlawful sharing: *Sabam v. Netlog*

In 2008, the CJEU ruled in the *Promusicae* case.¹⁵ An association of copyright-holders (Promusicae) requested that a telecom provider (Telefonica) provide them with the identity and physical addresses of subscribers they suspected of infringing their copyright via a peer-to-peer file-sharing service. Promusicae wished to initiate civil proceedings against the alleged infringers. The CJEU found (1) that MSs are not obliged to impose an obligation to communicate personal data such as those requested, but (2) that they may nevertheless impose such an obligation, taking into account the eCommerce Directive, the Copyright Directive, the IP Enforcement Directive, the ePrivacy Directive, and the Data Protection Directive, while also implementing the proportionality principle that informs the balancing of fundamental rights. Though this may sound rather vague, it is actually quite precise, pointing out that a complex set of interacting parameters, derived from the relevant directives, must be applied to the case at hand in a way that does not violate fundamental rights and principles of the legal order of the EU. The precision will come into its own when the national court performs the required balancing act.

In more recent case law, the CJEU has ruled against the imposition of monitoring and filtering obligations on ISPs, based on them violating the above quoted Article 15 of the eCommerce Directive. In *Sabam v. Netlog*,¹⁶ Sabam represented copyright-holders, authorizing usage of copyrighted material by third parties, whereas Netlog was an online social network that allowed file-sharing. Sabam had filed an injunction against Netlog to cease unlawful sharing, whereas Netlog argued that this would constitute a general obligation to monitor all its customers, as a preventative measure. The Belgium court submitted the following preliminary questions to the CJEU:

Do Directives 2001/29 and 2004/48, in conjunction with Directives 95/46, 2000/31 and 2002/58, construed in particular in the light of Articles 8 and 10 of the European Convention on the Protection of Human Rights and Fundamental Freedoms (...),

- permit Member States to authorise a national court, (...)

¹⁵ CJEU, 29 January 2008, C-275/06 (*Promusicae*).

¹⁶ CJEU, 16 February 2012, C-360/10 (*Sabam v. Netlog*).

- on the basis merely of a statutory provision stating that ‘[the national courts] may also issue an injunction against intermediaries whose services are used by a third party to infringe a copyright or related right’,
- to order a hosting service provider to introduce,
- for all its customers,
- *in abstracto* and as a preventive measure,
- at its own cost and for an unlimited period,
- *a system for filtering* most of the information which is stored on its servers
- in order to identify on its servers electronic files containing musical, cinematographic or audio-visual work
- in respect of which SABAM claims to hold rights,
- and subsequently to block the exchange of such files?

Taking into account the relevant directives, the CJEU decided in considerations 42–50, that as:

- IP rights are fundamental rights to property,
- a fair balance must be struck between property rights (Article 17 CFREU);
 - the freedom to conduct a business (hosting service) (Article 16 CFREU);
 - the right to protection of the personal data of the user (Article 8 CFREU);
 - the user’s freedom to receive or impart information and the freedom of information (Article 11 CFREU).

It then finds that such a balance precludes an injunction for Netlog to systematically, indiscriminately, preventively monitor and filter, exclusively at its expense and for an unlimited period. I leave it to the reader to carefully excavate the precise reasons for this categorically precise answer to the questions raised.

7.4.1.5 Injunctions to cease unlawful sharing: *Brein v. Ziggo*

In 2017, the CJEU decided the case of *Brein v. Ziggo*,¹⁷ about the lawfulness of a court order to ISPs to block The Pirate Bay (TPB). This case is about the Netherlands Foundation Brein, which represents copyright-holders, applying for a court order to ISPs Ziggo and XS4ALL, to block TPB, claiming that TPB violates copyright. The court of first instance has granted the application. The court of appeal has denied the application, because although uploading and downloading of illegally distributed content is illegal, empirical

¹⁷ CJEU, 14 June 2017, Case C-610/15 (*Stichting Brein v. Ziggo*).

evidence shows that blocking is not effective and therefore the measure is not proportional.

In its judgment,¹⁸ the Netherlands Supreme Court notes that the CJEU has decided in *UPC Telekabel Wien*,¹⁹ that even if a measure does not rule out violation this does not necessarily imply that it cannot be proportional; instead, the CJEU found that a measure may be proportional as long as the measure makes violation more difficult and provides a clear sign that users should not engage in illegal uploading and downloading. The Supreme Court concludes from this that the court of appeal applied a wrong criterion to test the proportionality requirement, its interpretation being too strict. The Supreme Court subsequently submits preliminary questions to the CJEU, concerning the question whether TPB itself should be considered as committing copyright infringements. This question is interpreted by the CJEU as follows:

18. By its first question, the referring court asks, in essence, whether the concept of ‘communication to the public’, within the meaning of Article 3(1) of Directive 2001/29, should be interpreted as covering, in circumstances such as those at issue in the main proceedings, the making available and management, on the internet, of a sharing platform which, by means of indexation of metadata relating to protected works and the provision of a search engine, allows users of that platform to locate those works and to share them in the context of a peer-to-peer network.

The Court responds:

24. It is clear from Article 3(1) of Directive 2001/29 that the concept of ‘communication to the public’ involves two cumulative criteria, namely an ‘act of communication’ of a work and the communication of that work to a ‘public’ (...).

26. (...) That user makes an act of communication when he intervenes, in full knowledge of the consequences of his action, to give his customers access to a protected work, particularly where, in the absence of that intervention, those customers would not be able to enjoy the broadcast work, or would be able to do so only with difficulty (...)

27. (...) the concept of the ‘public’ refers to an indeterminate number of potential viewers and implies, moreover, a fairly large number of people

¹⁸ Netherlands Supreme Court, 13 November 2015, ECLI:NL:HR:2015:3307 (*TPB*).

¹⁹ CJEU, 27 March 2014, C-314/12 (*UPC Telekabel Wien*).

28. (...) according to a settled line of case-law, in order to be categorised as a ‘communication to the public’, a protected work must be communicated using specific technical means, different from those previously used or, failing that, to a ‘new public’, that is to say, to a public that was not already taken into account by the copyright holders when they authorised the initial communication of their work to the public (...)

29. (...) the profit-making nature of a communication, within the meaning of Article 3(1) of Directive 2001/29, is not irrelevant (...).

39. In the light of the foregoing, the making available and management of an online sharing platform, such as that at issue in the main proceedings, must be considered to be an act of communication for the purposes of Article 3(1) of Directive 2001/29.

43. It follows that, by a communication such as that at issue in the main proceedings, protected works are indeed communicated to a ‘public’ within the meaning of Article 3(1) of Directive 2001/29.

45. (...) In any event, it is clear from the order for reference that the operators of the online sharing platform TPB could not be unaware that this platform provides access to works published without the consent of the rightholders, given that, as expressly highlighted by the referring court, a very large number of torrent files on the online sharing platform TPB relate to works published without the consent of the rightholders. In those circumstances, it must be held that there is communication to a ‘new public’ (...).

46. Furthermore, there can be no dispute that the making available and management of an online sharing platform, such as that at issue in the main proceedings, is carried out with the purpose of obtaining profit therefrom, it being clear from the observations submitted to the Court that that platform generates considerable advertising revenues.

47. Therefore, it must be held that the making available and management of an online sharing platform, such as that at issue in the main proceedings, constitutes a ‘communication to the public’, within the meaning of Article 3(1) of Directive 2001/29.

This results in the following decision:

On those grounds, the Court (Second Chamber) hereby rules:

The concept of ‘communication to the public’, within the meaning of Article 3(1) of Directive 2001/29/EC (...), must be interpreted as covering, in circumstances such

as those at issue in the main proceedings, the making available and management, on the internet, of a sharing platform which, by means of indexation of metadata relating to protected works and the provision of a search engine, allows users of that platform to locate those works and to share them in the context of a peer-to-peer network.

7.4.1.6 The update of the Copyright Directive

In 2019, an update has been enacted with an eye to ‘copyright and related rights in the Digital Single Market’, which for instance provides an exception for reproductions or extractions to the extent that they are used for text and data mining.²⁰ The update caused an uproar around two other articles, notably Articles 15 and 17.

Article 15 aims to protect newspaper publishers against sharing of their content by providing them with the exclusive right to authorize or prohibit reproduction and online publication, requiring content aggregators to compensate these publishers by way of a so-called ‘link-tax’. This is deemed by many to seriously harm the freedom of information. Article 17 aims to protect copyright-holders against illegal sharing of their works via peer-to-peer sharing platforms, by overruling the exemption of liability of intermediaries of Article 14 of the eCommerce Directive. This basically requires extensive filtering of content to check whether protected content is uploaded without right. This is deemed by many to strengthen the powers of already powerful big players because they can afford to invest huge amounts of money in filtering software, whereas small players may shrink from even entering the market, fearing either huge compliance costs or high risk to incur liability for unwittingly sharing illegal content. On top of that, many find the delegation of policing tasks to private companies highly problematic, as this may result in unintended censure of lawful content. The European Parliament, however, stated in a press statement that ‘[m]aking internet companies liable will enhance rights holders’ chances (notably musicians, performers and script authors, as well as news publishers and journalists) to secure fair licensing agreements, thereby obtaining fairer remuneration for the use of their works exploited digitally.’²¹

An in-depth discussion of the pros and cons of these articles is outside the scope of this book, but I encourage those interested to study the relevant articles of the directive and the relevant debates in the relevant scholarly literature.

²⁰ For details, see Articles 3 and 4 of Directive (EU) 2019/790.

²¹ News European Parliament, ‘European Parliament approves new copyright rules for the internet’, 26 March 2019, available at: <http://www.europarl.europa.eu/news/en/press-room/20190321IPR32110/european-parliament-approves-new-copyright-rules-for-the-internet>.

7.4.2 The Software Copyright Directive

The object of protection of the Software Copyright Directive is defined in Article 1, as a specific type of ‘literary works’ (within the meaning of the Berne Convention, see section 7.3), namely *computer programs*, including their preparatory design material, which concerns—as always in copyright—only the expression, not ideas or principles, including those which underlie its interfaces. Protection is only available if the program is original in the sense of being an author’s own intellectual creation.

Articles 3 and 2 determine that the *right-holder* is the author of the program, defined as the *creator of the program*. However, if created by an employee, the employer is entitled to exercise all economic rights. If a group of natural persons jointly created the program, the exclusive rights will be owned jointly.

Article 4 defines the *restricted acts* (scope of protection) with regard to a computer program due to the exclusive nature of a copyright: reproduction; translation, adaptation, arrangement, any other alteration; and distribution to the public (publication), including by way of rentals. First sale in the Community exhausts the right of distribution.

Article 5 exhaustively defines *exceptions* to the copyright (limitation of what acts can be restricted). Unless prohibited by contract, no authorization is required for the reproduction and translation of a program if necessary for a reasonable use of the program, with the exception of the making of a backup copy which cannot be prohibited by contract insofar as necessary for the use of the program. A person who has the right to use the program is allowed to observe, study, or test the functioning of a program to determine underlying ideas and principles, if done ‘while performing any of the acts of loading, displaying, running, transmitting or storing the program which he is entitled to do’.

Article 6 addresses *decompilation*, stipulating that reproduction and translation of code (as defined in Article 2) are allowed to achieve interoperability, if specified conditions are met (notably that such reproduction or translation is performed by the licensee or another person with a right to use a program; that the interoperability is not already readily available; that such reproduction or translation is confined to parts of the original program that are necessary for interoperability). Information obtained in order to achieve interoperability may not be further used for other purposes, nor be shared with others if not necessary to achieve interoperability, nor be used to create a similar program as the original one.

As the reader will note, I have not quoted the *original* articles of the directive, but paraphrased them and thus compressed their *underlying ideas*. This is

risky, because it is only the *expression* in the wording of the directive that has force of law. Compression, as any computer scientist knows, implies an act of interpretation. To analyse the relevant law in terms of legal effect and legal conditions, the precise wording must thereby be followed. Nevertheless, lawyers are used to such acts of compression and interpretation, in function of clarity or speed, whilst remaining deeply aware that for decision-making the original expression is authoritative (the reader is advised to consult the precise articulation of the relevant articles for themselves).

Even so, the original expression cannot speak for itself, requiring interpretation in the light of relevant facts to which it must be applied. As discussed above, this requires interaction between: (1) the text of the law that must be interpreted in the light of the facts; and (2) the facts of a case that must be interpreted in the light of the applicable law. To further illustrate this point, I will discuss two leading cases of the CJEU with regard to the interpretation of the Software Copyright Directive.

7.4.2.1 Exceptions to the exclusionary software copyright: *SAS v. WLP*

This case²² concerned the dividing lines between the exclusionary right in the object of protection and the exceptions relevant when seeking to uncover the underlying ideas and functionality of the program.

SAS developed a program for data analytics. Part of the program helps users to build their own modules which can then be used together with that same part of SAS's analytics program. World Programming Languages (WPL) sold a program that mimics the core part of SAS's program, thus creating an alternative for users of the SAS program. To protect its market share, SAS takes WPL to court for violation of its software copyright. This basically concerns the question of when reverse engineering is a violation of copyright and when it falls within the scope of the relevant exceptions.

The CJEU addressed three preliminary questions (as compressed and expressed in my own words):

1. Can the functionality of a computer program, the programming language, and format of data files be construed as a form of expression, and therefore be protected as copyright?

The CJEU responds by concluding that: all forms of expression that permit reproduction of the program, including source code and object code are *protected*; the graphic user interface, the functionality of the program, the

²² CJEU, 2 May 2012, C-406/10 (*SAS v. World Programming*).

programming language, and format of data files used by a program to exploit particular functions of the program are *not protected*.

2. What is the liability of a licensee—even if they are a competitor—who acts outside the scope of that licence to observe and study the functioning of a computer program in order to determine the ideas and principles behind that program?

The CJEU responds by concluding that: a licensee is entitled to observe, study, or test the functioning of a software program in order to determine the ideas and principles behind any and all elements of the program, as long as the licensee does not infringe copyright, for example, by using the source code or the object code. The Court finds that this means that the licensee is entitled to determine the ideas and principles while loading, displaying, running, transmitting, or storing the program. The Court also finds that copyright-owners cannot use the contractual nature of the software licence to stop licensees from performing acts necessary to observe, study, or test the functioning of the program, as long as these acts do not infringe the copyright in that program.

3. Does the reproduction in a program (or user manual for that program) of material described in a user manual for another (copyrighted) program constitute copyright infringement?

The CJEU concludes that: the reproduction of particular elements in a user manual for a computer program *may constitute a copyright infringement, if the material reproduced constitutes an expression of the author's intellectual creation*. Whereas keywords, syntax, commands, combinations of commands, options, defaults, and iterations consist of words, figures, or mathematical concepts *are not protected by copyright by themselves, they may be protected if combined in a manner that constitutes an intellectual creation*. Whether the choice, sequence, and combination of words, figures, or mathematical concepts is indeed an intellectual creation that constitutes a copyright in the user manual, will be *a matter of fact and degree*.

7.4.2.2 Exceptions to the exclusionary software copyright: *Microsoft*

This case²³ concerns the interpretation of Article 4.2 and Article 5.1 and 2 of the Software Copyright Directive. It came about ‘in the context of criminal proceedings brought by the (...) Department for the Prosecution of Economic and Financial Offences in Latvia against Mr Aleksandrs Ranks and Mr Jurijs Vasiļevičs, charged with the unlawful sale, as part of a criminal

²³ CJEU, 12 October 2016, C-166/15 (*Microsoft*).

organisation, of objects protected by copyright, (...) having sold, through an online marketplace, used copies of computer programs stored on non-original media' (paragraph 2 in the *Microsoft* case).

The Court reminds us that the term 'sale' in Article 4.2 (exhaustion after first sale) 'must be given a broad interpretation, encompasses all forms of marketing of a copy of a computer program characterised by the grant of a right to use that copy, for an unlimited period, in return for payment of a fee designed to enable the copyright holder to obtain a remuneration corresponding to the economic value of that copy' (paragraph 28).²⁴

The Court then decides that 'it follows from the Court's case-law that Article 4.2 (...), refers, without further specification, to the 'sale ... of a copy of a program' and thus makes no distinction according to the tangible or intangible form of the copy in question' (paragraph 35).²⁵ This implies that a lawfully downloaded copy must be considered equivalent with a copy stored on a DVD, as far as the exhaustion of first sale is concerned. Note the difference with the general Copyright Directive, which constrains this limitation to the first sale of tangible copies.²⁶

The facts of the case, however, concern the resale of a backup copy of the relevant software, because Mr Ranks and Mr Vasiļevičs no longer had access to the original copy. They argued that as they had a right to make a backup copy in order to use the original copy (Article 5.2), they could sell such a backup copy under the exception of exhaustion after first sale of Article 4.2. The CJEU, however, finds that 'a back-up copy of a computer program may be made and used only to meet the sole needs of the person having the right to use that program and that, accordingly, that person cannot—even though he may have damaged, destroyed or lost the original material medium—use that copy in order to resell that program to a third party' (paragraph 43). The CJEU therefore rules:

Article 4(a) and (c) and Article 5(1) and (2) of Council Directive 91/250/EEC of 14 May 1991 [now Directive 2009/24/EC] on the legal protection of computer programs must be interpreted as meaning that, although the initial acquirer of a copy of a computer program accompanied by an unlimited user licence is entitled to resell that copy and his licence to a new acquirer, he may not, however, in the case where the original material medium of the copy that was initially delivered to him has been damaged, destroyed or lost, provide his back-up copy of that program to that new acquirer without the authorisation of the rightholder.

²⁴ Referring to *UsedSoft* (n 2 above) at para. 49.

²⁵ *Ibid.* at para. 55.

²⁶ See recital 28 and article 4 of Copyright Directive 2001/29/EC, and, above, footnote 6.

7.5 Open Source and Free Access

In 1983, Richard Stallman initiated the GNU project, publishing the GNU Manifesto, where he explains:

GNU, which stands for Gnu's Not Unix, is the name for the complete Unix-compatible software system which I am writing so that I can give it away free to everyone who can use it. Several other volunteers are helping me. Contributions of time, money, programs and equipment are greatly needed.

In 1985, he founded the Free Software Foundation (FSF), meant to develop and share software based on GPLs.

In 1991, Linus Torvalds developed the LINUX kernel, aiming to enable software and hardware to interact. Together with the GNU software developed in the context of the GNU project, LINUX forms an operating system, though by now this is also achieved with other types of software that is not necessarily part of the GNU project. LINUX is free to use, and everyone has the freedom to contribute to its development.

In 1998, the Open Source Initiative (OSI) was founded, referring not merely to the freedom to use software, but including the necessity to disclose the source code.

The FSF defines free software in terms of four freedoms:²⁷

1. The freedom to run the program as you wish, for any purpose (freedom 0).
2. The freedom to study how the program works, and change it so it does your computing as you wish (freedom 1). Access to the source code is a precondition for this.
3. The freedom to redistribute copies so you can help others (freedom 2).
4. The freedom to distribute copies of your modified versions to others (freedom 3). By doing this you can give the whole community a chance to benefit from your changes. Access to the source code is a precondition for this.

To ensure that software remains 'free software' in the above sense, a variety of software licences has been developed that contractually require that the software is further developed and shared 'freely' in the above sense. This can be

²⁷ GNU website: <https://www.gnu.org/philosophy/free-sw.html#f1>. The website clarifies in footnote [1]: 'The reason they are numbered 0, 1, 2 and 3 is historical. Around 1990 there were three freedoms, numbered 1, 2 and 3. Then we realized that the freedom to run the program needed to be mentioned explicitly. It was clearly more basic than the other three, so it properly should precede them. Rather than renumber the others, we made it freedom 0.'

done in an absolute way, meaning that: the freedom goes viral; that no restrictions are possible by subsequent users, for subsequent versions; that each derivative work is contaminated by the same requirements by means of the same licence (this is referred to as ‘copyleft’). It can also be done in a less absolute manner, meaning that: such radical viral effect is not necessary, for instance allowing subsequent versions to be part of a proprietary licence (this is referred to as ‘non-copyleft’).

In 2001, Lawrence Lessig initiated the Creative Commons (cc), transposing the idea of Open Source to other creations (non-software). The Creative Commons have developed a set of different licences, enabling a more granular scale of control over subsequent versions of the same creative expression.²⁸

In Figures 7.1 and 7.2 the reader can see the kind of—limited—granularity this provides, especially when combining different conditions.

Attribution (by)

All CC licenses require that others who use your work in any way must give you credit the way you request, but not in a way that suggests you endorse them or their use. If they want to use your work without giving you credit or for endorsement purposes, they must get your permission first.

ShareAlike (sa)

You let others copy, distribute, display, perform, and modify your work, as long as they distribute any modified work on the same terms. If they want to distribute modified works under other terms, they must get your permission first.

NonCommercial (nc)

You let others copy, distribute, display, perform, and (unless you have chosen NoDerivatives) modify and use your work for any purpose other than commercially unless they get your permission first.

NoDerivatives (nd)

You let others copy, distribute, display and perform only original copies of your work. If they want to modify your work, they must get your permission first.

Figure 7.1 Creative Commons Licence Conditions

²⁸ <https://creativecommons.org/share-your-work/licensing-types-examples/>.


Icon	Description	Acronym	Free Cultural Works	Remix culture	Commercial use
	Freeing content globally without restrictions	CC0	Yes	Yes	Yes
	Attribution alone	BY	Yes	Yes	Yes
	Attribution + ShareAlike	BY-SA	Yes	Yes	Yes
	Attribution + Noncommercial	BY-NC	No	Yes	No
	Attribution + NoDerivatives	BY-ND	No	No	Yes
	Attribution + Noncommercial + ShareAlike	BY-NC-SA	No	Yes	No
	Attribution + Noncommercial + NoDerivatives	BY-NC-ND	No	No	No

Figure 7.2 Types of licences

By now, similar open access models have been developed in the realm of patent law, for instance by building publicly available databases (scientific research, e.g. human genome project) to nourish the public domain, and as a defence strategy, since public release prevents patenting. This can be combined with copyleft prohibition of downstream restrictions, for instance with HAPMAP (only regards the data, not derived applications), BiOS (application to the patented invention, though improvements can be patented).

With special regard to developing countries, an equitable access licence and a neglected disease licence have been proposed. Finally, we can observe further consolidation in the 2001 Budapest Open Science Initiative, the 2003 Berlin Declaration on Open Access to Knowledge in the Sciences and the Humanities, and the 2007 Science Commons.

The common core of all these varieties of open source, free software, and open access models is:

1. the assertion of IP right;
2. the reverse use of exclusivity; and
3. the absence of discrimination.

It is crucial to remember that Creative Commons licences make no sense without the underlying property right in the work that is developed and shared.

References

On the history of copyright

Rose, Mark. 1995. *Authors and Owners: The Invention of Copyright*. Reprint edition. Cambridge, MA: Harvard University Press.

On the nature of IP rights

Cohen, Julie E. 2014. 'What Kind of Property is Intellectual Property?' *Houston Law Review*. Symposium 2014, 52 (12): 691–707.

Introduction to IP law

Cook, Trevor. 2010. *EU Intellectual Property Law*. Oxford and New York: Oxford University Press.

Hunter, Dan. 2012. *The Oxford Introductions to U.S. Law: Intellectual Property*. New York: Oxford University Press.

The relationship between IP and Creative Commons

Boyle, James. 2009. 'What Intellectual Property Law Should Learn from Software'. *Communications of the ACM* (52) (September), 71–76. <http://www.thepublicdomain.org/2009/08/26/what-intellectual-property-law-should-learn-from-software/>.

Dusollier, Severine. 2007. 'Sharing Access to Intellectual Property through Private Ordering'. *Chicago-Kent Law Review* 82; 1391–435. http://works.bepress.com/severine_dusollier/14/.

Lessig, Lawrence. 2006. *Code Version 2.0*. New York: Basic Books, chapter 10.

Private Law Liability for Faulty ICT

What if a new version of an operating system (OS) is launched, enabling users to upgrade automatically or manually? Do such upgrades have legal effect?

- Maybe those who upgrade manually run security risks, including breakdown of their application, if they fail to implement the upgrade in good time.
- Maybe previous versions of the OS will not be supported after some time (no updates), meaning that those running the OS on hardware that does not support the upgrade will be left unprotected.
- Maybe those who subscribed to automated updates on the new version of the OS inadvertently install spyware that causes harm to their private life, business interest, or employee status.

What if one's smart fridge communicates with a host of providers (from the hardware manufacturer to the providers of the OS and various applications, such as those of online groceries or health insurance providers that monitor eating habits)? Does such communication have legal effect?

- Maybe the fridge is confronted with power cuts due to issues in the smart energy grid that results in electrocution, because this scenario has not been foreseen by its adaptive software.
- Maybe the fridge is run based on a smart contract, implemented via a blockchain that disconnects the fridge due to a default in payment, causing a short circuit.
- Maybe the fridge starts ordering the same food from a number of different groceries due to a bug in its system, whereas these contracts are automatically executed without recourse to nullification.

The reader can easily imagine other instances where ICT—whether adaptive or self-executing—causes physical, material, economic, or emotional harm. For instance, what if one misses an important appointment due to the washing machine catching fire (material damage to the machine, the bathroom), which causes one to default on a contract that results in loss of income (economic damage), or what if one witnesses fearsome bodily harm or death

of a close relative due to the fire which results in a post-traumatic stress syndrome (emotional damage)? Maybe the fact that one's personal data have been leaked by an insurance company in a major data breach causes enduring anxiety about who may have accessed, sold, or otherwise shared the data.

The question of whether having caused such damage could have legal effect is a matter of tort law. Legal conditions for such legal effect demand that these harms can be attributed to, for example, the manufacturer, the operating system (OS) provider, the retailer, the insurance company where a breach occurred, the helpdesk provider that gave the wrong advice, or the firm that leased the car, the washing machine, or the fridge (as this firm may have changed the manufacturer's default settings, thus causing the harm).

Maybe, on the other hand, the washing machine simply does not function as well as before, ever since an update has been installed. Maybe the brakes of a connected car are in turbulence due to a bug in the OS. This raises the question of whether one could sue the seller based on non-conformity of the product or service with what one could reasonably expect, considering its function and the price, or on the basis of a defect.

In this chapter I will focus on third party liability, or tort law, as an important example of how private law liability may step in to deter developers, manufacturers, sellers, and users of ICT from developing, selling, or using faulty ICT.

8.1 Back to Basics

Before moving head on into third party liability we first revisit the basics presented in the first part of this book.

8.1.1 Chapter 3: private law distinctions

In private law we discriminate between *absolute and relative rights*, where absolute rights play out in the relationships between a legal subject and all other legal subjects (within a jurisdiction) with regard to a specific object (a movable, real estate, or an immaterial good such as a work or an invention, or even with regard to a receivable).¹

¹ A receivable is a relative right to receive a payment. This is considered to be an asset that can be traded or sold.

In case of an absolute right, all others must refrain from interfering with the object. Relative rights play out between designated legal subjects, such as the parties to a contract or a tortfeasor and their victim. Liability based on tort is called third-party liability, because it is not based on the direct relationship between the parties to a contract, but involves a third party. In some jurisdictions it is possible to issue a tort action against one's contracting party. This means that one does not base the action on breach of contract, but on the other party being liable for damage on grounds of tort.

As discussed in Chapter 3, the purpose of private law can be summarized under the headings of (1) respecting individual autonomy; (2) ensuring fairness, such as compensation of inequality that would diminish individual autonomy, which may require a party to, for example, inform the other party or to shift the burden of proof to the party with access to relevant evidence; and (3) the societal trust that is pivotal for the functioning of economic markets. Private law is restricted by constitutional limitations (e.g. government may, under strict conditions, dispossess the owner of real estate in the general interest), by international human rights law (e.g. horizontal effect of privacy), and by administrative law (e.g. requiring a permit to renovate one's own property).

Private law contains more *default law*, especially in the domain of contract law, where the freedom to contract often implies that contracting parties may deviate from the legal provisions that would otherwise rule their contract. Property law contains more *mandatory law*, due to its third-party effects (property law affects all others as it concerns absolute rights). Due to the legality principle, public law contains mostly mandatory law (as legal powers of government bodies should clarify what citizens can expect).

The legality principle also plays a major role in the criminal law. The set of all unlawful cyber conduct contains, for example, cyber torts, violations of cyber-related contracts, and violations of cyber-related administrative law. Only a small part of this set concerns cybercrime, because only unlawful conduct that has been explicitly criminalized constitutes a crime.

Finally, let's once again consider the notions of a legal subject and a legal object.

1. A legal subject (a natural person or legal person) is an entity capable of acting in law, bearing legal rights and legal obligations in relation to other legal subjects.

2. A legal object (a good: intellectual property rights; real estate; tangibles; other rights and obligations) is an entity that is the object of legal relationships between legal subjects.

In the case of a tort, the legal subjects are the tortfeasor and the victim, while the legal object is a prohibition to engage in tortuous conduct (in the case of an injunction) and/or an obligation to pay damages and the right to be compensated for the damage one suffered.

8.1.2 Chapter 4: international and supranational law

International private law (IPL) concerns ‘the law of conflicts’ that determines applicable law and the jurisdiction of national courts in cases where different jurisdictions may be applicable both regarding the substance (which law is applicable?) and regarding the competence of a court (which national courts have the power to admit a case?). In the end, IPL is national law, since national law decides whether its courts have competence and what law they should apply. As this leads to conflicts whenever different states decide differently on the same case, international treaties have been concluded to prevent overlapping jurisdiction or conflicting applicable content.

There is no supranational private law, despite numerous attempts to agree on a ‘European private law’, so third-party liability for faulty ICT cannot be based on EU tort law. In the context of the goal of creating and sustaining the internal EU market there are many reasons for such a ‘common’ private law, as it would increase legal certainty for companies providing products and services across national borders, achieving at least minimal harmonization that would also protect EU consumers and small companies, while preventing and reducing unbalanced competition (market entry) and administrative burdens, (which vary depending on requirements stemming from national tort law).

There is, however, a set of relevant EU directives that requires harmonization on issues that overlap with tort law,² such as the Product Liability Directive,³ the Unfair

² The same goes for EU directives that require MSs to implement private law regarding contract law, such as Directive 93/13/EC on unfair terms in consumer transactions (what should be blacklisted, what should be greylisted), Directive 97/7/EC on distant selling and Directive 99/44/EC on the sale of consumer goods.

³ Directive 85/374/EC on liability for defective products, raising the question of what is a product and who qualifies as a producer.

Commercial Practices Directive,⁴ the eCommerce Directive (see section 7.4.1.3 above),⁵ and the ePrivacy Directive (with a potentially new liability regime under the upcoming ePrivacy Regulation).⁶ Being directives, the harmonization is somewhat limited by the fact that member states (MSs) have to implement the directives into their national legal framework, instead of having to adhere to one and the same text. Nevertheless, where such directives require MSs to enable tort liability (or provide exemptions), the Court of Justice of the European Union (CJEU) usually finds that such requirements must be understood in an autonomous manner that enables a consistent interpretation throughout the Union.⁷ Article 82 of the General Data Protection Regulation (GDPR) may become an interesting example of such an ‘autonomous’ interpretation (see above section 5.5.2.11 and below section 8.1.3).

8.1.3 Chapter 5: data protection law

The GDPR has a specific chapter that is dedicated to the enforcement of the regulation (see also above section 5.5.2.11). This chapter contains the following articles: Article 77 ‘Right to lodge a complaint with a supervisory authority’; Article 78 ‘Right to an effective judicial remedy against a supervisory authority’; Article 79 ‘Right to an effective judicial remedy against a controller or processor’; Article 80 ‘Representation of data subjects’; Article 82 ‘Right to compensation and liability’; Article 83 ‘General conditions for imposing administrative fines’; Article 83.1 ‘effective, proportionate and dissuasive’, Article 83.4 ‘maximum 2% global turnover’, Article 83.5 ‘maximum 4% global turnover’; Article 84 ‘Penalties’, especially for infringements not subject to the fines of Article 83, and those penalties should again be effective, proportionate, and dissuasive.

So far, most of the attention has focused on the fines. The chapter contains a very smart set of private law remedies, however, that may provide highly effective incentives to companies processing personal data.

⁴ Directive 2005/29/EC of the European Parliament and of the Council concerning unfair business-to-consumer commercial practices in the internal market.

⁵ This could concern ‘fake news’, botfarms etc.

⁶ Article 22, para. 1 of the Commission proposal reads: ‘Any end-user of electronic communications services who has suffered material or non-material damage as a result of an infringement of this Regulation shall have the right to receive compensation from the infringer for the damage suffered, unless the infringer proves that it is not in any way responsible.’ EP Amendment 50 adds: ‘Article 82 of Regulation (EU) No 2016/679 shall apply.’

⁷ CJEU, 12 March 2002, C-168/00 (*Leitner*); CJEU, 25 October 2005, C-350/03, (*Schulte*); and CJEU, 2 June 2005, C-229/04 (*Crailsheimer Volksbank*).

Article 79—Right to an effective judicial remedy against a controller or processor

1. Without prejudice to any available administrative or non-judicial remedy, including the right to lodge a complaint with a supervisory authority pursuant to Article 77, each data subject shall have
 - the right to an effective judicial remedy
 - where he or she considers that
 - his or her rights under this Regulation have been infringed
 - as a result of the processing of his or her personal data
 - in non-compliance with this Regulation.

This article basically stipulates that data subjects should be able to lodge an injunction against a controller they believe to be unlawfully processing their personal data. As the remedy must be ‘effective’, we may expect court orders to be reinforced with penalty payments in case of non-compliance.

As filing a court case is neither easy nor obvious for individual data subjects, Article 80 provides important possibilities for collective action, despite the fact that there is no consensus on a dedicated directive on Union-wide collective action (notably not for compensation of damages).

Article 80—Representation of data subjects

1. The data subject shall have the right to mandate
 - a not-for-profit body, organisation or association
 - which has been properly constituted in accordance with the law of a Member State,
 - has statutory objectives which are in the public interest,
 - and is active in the field of the protection of data subjects’ rights and freedoms
 - with regard to the protection of their personal data
 - to lodge the complaint on his or her behalf,
 - to exercise the rights referred to in Articles 77, 78 and 79 on his or her behalf, and
 - to exercise the right to receive compensation referred to in Article 82 on his or her behalf where provided for by Member State law.

This is very interesting because—whereas this leaves collective action regarding compensation up to the MSs—MSs will have to allow data subjects to mandate their right to file an injunction to prohibit unlawful processing based on Article 79 to a relevant not-for-profit body. The second paragraph of Article 80 also leaves up to the MSs the possibility to enable a relevant not-for-profit body to start such actions on their own behalf.

Article 82, finally, requires that MSs create private law liability for unlawful processing that causes harm or damage:

Article 82—Right to compensation and liability

1. Any person who has suffered material or non-material damage as a result of an infringement of this Regulation shall have
 - the right to receive compensation from the controller or processor for the damage suffered.
2. Any controller involved in processing shall be liable
 - for the damage caused by processing which infringes this Regulation.
 A processor shall be liable for the damage caused by processing only
 - where it has not complied with obligations of this Regulation
 - specifically directed to processors or
 - where it has acted outside or contrary to lawful instructions of the controller.
3. A controller or processor shall be exempt from liability under paragraph 2 if
 - it proves that it is not in any way responsible for the event giving rise to the damage.

Paragraphs 4 and 5 add several liability for joint controllers and the distribution of liability between processors and controllers.

As we will see in the next sections, tort liability comes in different shapes and versions. The reasonably granular stipulations of private law liability under the GDPR will probably contribute to legal certainty within the EU market regarding the legal effect of unlawful processing, thus avoiding different liabilities depending on different regimes of private law in the MSs.

8.2 Tort Law in Europe

In continental Europe, and in the parts of Africa, Latin America, and Asia that were influenced by its legal systems, private law has been codified by the legislator, such as the French *Code Civil* or the German *Bürgerliches Gesetzbuch* or the Netherlands *Burgerlijk Wetboek*. Such legal systems are usually referred to as the ‘civil law’ tradition. In Britain, the United States, Canada, Australia, and India, private law is part of the ‘common law’, which is based in ‘precedent’ or case law, rather than codification. This may lead to the conclusion that in civil law traditions code is all that matters, whereas in common law all depends on adherence to previous case law or precedent. Today, this is no longer the case. Whereas civil law takes its clue from legislation, the interpretation of the code

requires keen attention to prior case law; whereas common law takes its clue from prior case law, its interpretation requires keen attention to implied rules and principles that involve a similar systemization as aimed for by way of codification. Besides that, in common law jurisdictions, myriad statutory law has been enacted.

In this section, I will briefly revisit the main legal conditions that must be fulfilled to speak of a tortuous act (see also section 3.2.3 above). I will take into account the various civil and common law jurisdictions that ‘make up’ Europe, because as the United Kingdom has left the EU, economic intercourse within and between the United Kingdom and the EU will benefit from mutual recognition and proper understanding of the main pillars of tort law. In the light of remote access and remote control, enabled by hyperconnectivity and computational power, tort law will have to accommodate liability whenever a tort action has effects outside the jurisdiction where such action was initiated.

I will briefly discuss the requirements of damage, causation, fault liability, and strict liability, ending with questions around compensation and deterrence as the overarching goals of tort law.

Damage is the first requirement for a successful tort action, insofar as one wishes to obtain compensation. Such damage may refer to economic loss, personal injury, or a violation of personality rights; damages may be claimed for pain and suffering, for the violation of one’s dignity, and for the death of a beloved person.

So-called ‘wrongful life’ claims suggest that damage may even be established where a severely disabled person is born due to the violation of a duty of care by a healthcare institution that failed to notify the prospective parents of an increased risk of such a disability, thus preventing them from deciding to have an abortion.

Causation is the second requirement, since the damage must have been caused by the incriminated tortuous act to qualify for compensation. Usually, establishing causation refers to the so-called ‘*conditio sine qua non*’, which means that without the relevant act the damage would not have occurred.

This, however, refers to any action involved in the chain of events that led to the damage. The decision by the grandparent of the alleged tortfeasor to move to another country where they met their spouse-to-be is also a

conditio sine qua non, but will not be taken into account. To narrow down the ‘relevant cause’ we need a normative understanding of causation (which obviously has nothing to do with a ‘subjective’ understanding). Courts will take into account the remoteness of the damage, based on doctrines around ‘proximate cause’ (which seeks the nearest relevant cause rather than some remote forerunner), ‘adequate cause’ (which requires that the relevant action seriously increased the probability that the relevant damage would occur), or—more abstractly—‘reasonable attribution’ (which determines whether it is reasonable to qualify the relevant action as having caused the damage). Whatever theory is employed, to have the legal effect of a tort, a certain and direct causal connection is preferable to an uncertain and/or indirect connection. Beyond a certain threshold, causation will not be attributed.

The reduction of the space opened by the *conditio sine qua non* criterion is often achieved by taking into account the foreseeability of the damage. For instance, Dutch case law requires car drivers to foresee that other users of a public road will not comply with traffic rules. This means that cars may have to anticipate cyclists without light after dark. This clearly brings out the normative aspect of the causality attribution because, for instance, pedestrians crossing a zebra crossing need not anticipate speeding cars. Also, the wish to protect victims in personal injury cases and the blameworthiness of the tortfeasor can play a role in the attribution of causality. For instance, in the case of harm caused by asbestos or specific medication (e.g. DES).⁸ Often, it is not possible to identify which potential tortfeasors actually ‘caused’ what individual harm, for example, because the victim worked for several companies that used asbestos or the victim cannot prove which brand of medication was subscribed. Solutions to such problems are, for example, the imputation of several liability, or liability in proportion to the market share, often in combination with a reversal of the burden of proof, which brings us to the difference between different types of liability.

Liability regimes can be distinguished in terms of fault liability or strict liability, with some shades of grey in between.

Fault liability is based on the maxim that each victim bears their own damage, unless a special reason applies to shift the burden to another legal subject

⁸ DES was a drug given to pregnant women to prevent miscarriages. It was prescribed long after it was shown to be ineffective and gave rise to a number of law suits, class actions, and doctrinal debates within law, see e.g. <https://diethylstilbestrol.co.uk/sitemap/>.

who caused the damage. Such special reasons may be: (1) fault, which assumes intentional wrongdoing, or (2) negligence, which assumes a failure to exercise reasonable care.

Negligence is objectified by referring to the care that a reasonable person would or should have taken. Some diehard computationalists believe this can be caught in a formula. In *US v. Carroll Towing Co.*,⁹ the famous US Judge Learned Hand developed the following formula:

Since there are occasions when every vessel will break from her moorings, and since, if she does, she becomes a menace to those about her; the owner's duty, as in other similar situations, to provide against resulting injuries is a function of three variables: (1) The probability that she will break away; (2) the gravity of the resulting injury, if she does; (3) the burden of adequate precautions. Possibly it serves to bring this notion into relief to state it in algebraic terms: if the probability of an accident is called P, the injury L and the burden of precautions B, liability depends on whether B is less than L multiplied by P, i.e. whether $B < PL$.

As you may guess, a specific branch of Law and Economics (the so-called Chicago School) has picked up on this to develop intriguing theories on the utility of tort law as a means to prevent tortuous conduct. We will return to this issue under 'compensation and deterrence' at the end of this section.

An important extension of fault liability regards vicarious liability, which attributes the liability for tortuous conduct of one person to another person, often a legal person. For instance, the employer may be liable for tortuous conduct of their employees, insofar as the damage was caused in the normal course of the business.

Strict liability diverts from the baseline that each victim should bear their own damage. Here, damage is attributed without having to prove fault or negligence on the part of tortfeasor. This exception is often applied to a legal person that profits from the danger they create, considering that they are able to ensure against liability.

One can think of strict liability for:

- inherently dangerous people, goods, or activities (with so-called 'uncontrollable energy'), which are nevertheless not prohibited. For instance,

⁹ 159 F.2d 169 (2d. Cir. 1947) (*United States v. Carroll Towing Co.*).

strict liability of parents for the acts of their children; of car drivers that cause an accident involving a pedestrian; of pet-owners for their animals; or of the employer for work done at a construction site; and, perhaps, strict liability for the seller or user of inherently dangerous products and services with applied artificial intelligence (AI);¹⁰

- products or things that are not inherently dangerous but turn out defective for the purpose they were designed for (defective products, including products or services with applied AI).

Remedies in tort law can be distinguished as providing *compensation or deterrence* (or both). Tort law basically requires that people act as reasonable persons. Attributing liability for failure to do so enables one to shift the damage they cause from the victim to the tortfeasor, or at least to require monetary compensation. At the same time, it incentivizes potential tortfeasors to abstain from conduct that may cause damage.

As we have seen, sometimes tort law is used to compensate victims of damage caused by dangerous activities that society finds legitimate, such as driving a car. This implies that tort liability should not be confused with punishment and does not necessarily imply wrongfulness. It should therefore be distinguished from criminal law liability, but also from social security or private insurance on the side of the victim which may both compensate victims for harm and damage but will not have any deterrent effect on potential tortfeasors (who may then feel they can get away with dangerous conduct and thus externalize the costs of their decisions).

8.3 Third-Party Liability for Unlawful Processing and Other Cyber Torts

Third-party liability is defined as liability in the absence of a contract, where the victim and the tortfeasor do not have a direct relationship, which may, for example, cause difficulty in identifying the tortfeasor. The distance between victim and third party may be Euclidian (geographic) or otherwise, for instance due to the kind of network effects that ‘cyber’ applications generate. The emergence of cybercrime and the six relevant differences we identified compared with traditional crime, also apply to cyber torts: the differences in

¹⁰ With ‘user’ I do not refer to the end-user, but e.g. to a service provider that employs applied AI.

distance, scale, speed, distribution, invisibility, and visibility, brought about by the underlying automation and hyperconnectivity of networked computational systems (see section 6.1.2).

Some of the issues of third-party liability also relate to the political economy of big tech monopolies, often analysed under the heading of the ‘platform economy’. The exemption of liability for internet service providers (ISPs) (Articles 12–14 eCommerce Directive, discussed above at section 7.4.1.3) has resulted in difficulties for the allocation of responsibility in case of copyright infringements, child pornography, and ID theft that are enabled and mediated by ISPs. In the case of *Brein v. Ziggo* (regarding a court order to block The Pirate Bay (TPB), see above section 7.4.1.5), the CJEU decided that TPB—who claimed to be a mere intermediary—was itself infringing copyright. The CJEU thus allocated third-party liability to the ISP. Some may find this an infringement of the freedom of expression, as it requires ‘mere conduit’ ISPs to block ‘hosting’ ISPs, thus restricting the freedom of information of the users of the ‘hosting’ ISP. Similar arguments have been made regarding Article 13 of the upgrade of the Copyright Directive (see above section 7.4.1.6).

An important issue in the domain of third-party liability is that individual attempts to sue big players may not be effective in sustaining the societal trust that private law aims to achieve. This may be due to the fact that no concrete injury can be identified, that the costs of legal action overrule the benefits of compensation, or the simple fact that small players may not have the understanding, the time or the money to figure out how to assert their rights. One way to solve this problem is collective action, for instance by allowing people to mandate their claims to relevant not-for-profit associations, or by allowing a relevant not-for-profit to sue big players in their own name.

Such collective action may be very effective, especially in the case of (1) an injunction to arrest unlawful conduct, enforced with penalty payments for non-compliance, and in the case of (2) requesting a modest amount of compensation for a massive number of victims. As discussed in section 5.5.2.11, Article 80 GDPR offers new roads into an effective third-party liability regime, geared to prevent the relevant conduct by way of collective action. Though Article 80 does not require MSs to enable collective action to sue for damages, it does require them to enable collective action to stop unlawful processing.

8.3.1 Privacy harms

Finally, let me briefly discuss two examples of case law regarding ‘privacy torts’ within common law jurisdictions, which—as noted above—have a granular ‘law of torts’ rather than a general ‘tort law’ (as in civil law jurisdictions).

8.3.1.1 Canadian ‘tort of intrusion upon seclusion’

In *Jones v. Tsige*, the Canadian Court of Appeal for Ontario decided for the first time on a ‘tort of intrusion upon seclusion.’ The facts of the case are as follows:¹¹

In July 2009, the appellant, Sandra Jones, discovered that the respondent, Winnie Tsige, had been surreptitiously looking at Jones’ banking records. Tsige and Jones did not know each other despite the fact that they both worked for the same bank and Tsige had formed a common-law relationship with Jones’ former husband. As a bank employee, Tsige had full access to Jones’ banking information and, contrary to the bank’s policy, looked into Jones’ banking records at least 174 times over a period of four years.

The case is illuminating as it aims to establish whether the common law recognizes this type of privacy tort, based on an extensive investigation into common law jurisdictions (including Canada, the United States, and the United Kingdom). The Court argues that technological developments have indeed resulted in the need to recognize such a tort under common law.

They find that the legal effect of something qualifying as a ‘tort of intrusion upon seclusion’ depends on the following three legal conditions:¹²

1. the defendant’s conduct must be intentional (which includes recklessness);
2. the defendant must have invaded, without lawful justification, the plaintiff’s private affairs or concerns;
3. a reasonable person would regard the invasion as highly offensive, causing distress, humiliation, or anguish.

The ‘reasonable person’ test should prevent claims that are based on plaintiff’s subjective sensitivities or unusual privacy concerns.¹³

¹¹ 2012 ONCA 32 (*Jones v. Tsige*) at [2].

¹² *Ibid.* at [70] and [71].

¹³ *Ibid.* at [72].

The court also states that:¹⁴

Proof of harm to a recognized economic interest is not an element of the cause of action. (...) I believe it important to emphasize that given the intangible nature of the interest protected, damages for intrusion upon seclusion will ordinarily be measured by a modest conventional sum.

So, on the one hand the Court is willing to accept privacy harms that do not concern an economic interest, on the other hand the Court believes that the intangible nature of the harm implies compensation by way of ‘a modest conventional sum’. The compensation should remain symbolic, perhaps to highlight that such intangible harm resists monetisation.

8.3.1.2 UK ‘tort of misuse of private information’

In the case of *Murray v. Express Newspapers plc and another*,¹⁵ photographs were secretly taken (with a long-focus lens) of the young son of J.K. Rowling in a buggy, with his parents walking down a street. They were taken by a photographic agency, to be sold to interested parties, such as in this case the publisher of *The Sunday Express Magazine*, which published one of the pictures. The question whether this may constitute the tort of ‘misuse of private information’ was answered in reference to the following legal conditions:

- the plaintiff convincingly argues that they had a ‘reasonable expectation of privacy’ in the information; and
- the defendant cannot convincingly argue that a relevant justification applies, for instance claiming an overriding ‘public interest’ in publication.

The Court of Appeal extensively investigated the case law of the European Court of Human Rights (ECtHR) and the UK Data Protection Act (implementing the—then—applicable EU Data Protection Directive), to test whether the plaintiff could reasonably argue to have a ‘reasonable expectation of privacy’, and whether—if so—the proportionality test regarding justification could overrule such expectation.

It also discussed and rejected the verdict of the court of first instance with regard to its assessment of whether the plaintiff could substantiate damage.¹⁶

¹⁴ Ibid. at [71].

¹⁵ [2008] EWCA Civ 446 (*Murray v. Express Newspapers plc and another*).

¹⁶ Ibid. at 488 C.

Damage is not restricted to physical damage but includes pecuniary loss. An award is compensatory and includes the loss of the chance to sell the confidential information in question (...).

The conclusion here must be that within the context of the common law, old and new types of privacy torts are developing, due to the changing technological landscape.

8.3.2 Cyber torts?

Under the heading of cybercrime, we discussed the difference that makes a difference between cybercrimes and traditional crimes. As mentioned above, similar differences apply to the idea of cyber torts. We can, for instance, think of damage caused by malware, illegal access, ID fraud, domain hacking, by bullying, stalking, defaming, humiliating, grooming, by blocking access or availability, and by time-consuming and irritating communications such as spam. Privacy harm informed by hyperconnected computational systems could easily fall within the scope of cyber torts.

Types of torts could include: data breaches, unlawful processing of personal data, but also third-party liability for damage caused by non-conformity in the sale of goods or services, reputation damage, and safety hazards.

Types of damage could include: compensatory or punitive damages; direct and consequential damages; loss of earnings, loss of earning capacity, or loss of profit; material and immaterial damages; and present or future injury.

The examples given in the beginning of this chapter, highlighting damage caused by connected cars, smart fridges and intelligent washing machines, on the cusp of robotics, cloud robotics and the internet of things (iot), clearly raise a number of questions about the scope of the duty of care, the role of foreseeability when defining intent in the context of machine learning applications, issues of distributed causality in the case of integrated software and hardware components, the responsibility of the end-user for eventual consequential damage to others, and the extent to which unlawful processing of personal data in itself could be qualified as immaterial damage under EU data protection law, irrespective of the subjective experience of a data subject.

I expect that private law liability, together with data protection law, competition law, and consumer protection, will take the lead in reconfiguring the legal landscape of the onlife world. This should contribute to more adaptive legal protection and a better distribution of checks and balances between technology developers, manufacturers, retail, service providers, and end-users.

References

A short cross-jurisdictional introduction to tort law

Smits, Jan M. 2016. *Advanced Introduction to Private Law*. Cheltenham and Northampton, MA: Edward Elgar.

On privacy harms

Calo, Ryan. 2014. 'Privacy Harm Exceptionalism'. *Colorado Technology Law Journal* 12 (2): 361–64.

Solove, Daniel J., and Danielle K. Citron. 2018. 'Risk and Anxiety. A Theory of Data-Breach Harms'. *Texas Law Review* 96 (4): 737–86.

Discussions of cyber torts

Koch, Bernhard A. 2014. 'Cyber Torts: Something Virtually New?' *Journal of European Tort Law* 5 (2): 133–64. <https://doi.org/10.1515/jetl-2014-0009>.

Rustad, Michael L., and Thomas H. Koenig. 2005. 'The Tort of Negligent Enablement of Cybercrime'. *Berkeley Technology Law Journal* 20 (4): 1553. <https://doi.org/doi:10.15779/Z38JX0S>.

Rustad, Michael L., and Thomas H. Koenig. 2005. 'Harmonizing Cybertort Law for Europe and America'. *Journal of High Technology Law* 5: 13.

PART III

FRONTIERS OF LAW IN AN ONLIFE WORLD

In the third part of this book we will investigate two forward looking perspectives in and of law. These perspectives concern law's relationship with an environment that is increasingly data- and code-driven, where the difference between online and offline becomes ever more artificial. A difference to be *made*, instead of taking it for granted as *given*. From self-driving cars to online micro-targeting and from remote healthcare to e-learning, our physical, online, and institutional environment are integrating into an assembly of hard- and softwired decision-systems that interact and behave in myriad potentially unpredictable ways. Whether due to bugs, emergent properties or unforeseen dynamics. In Chapter 9 we will consider the salience of introducing legal personhood for some of the computational systems that run our world, thus also inquiring into the nature of legal subjectivity. In Chapter 10 we will examine the idea of 'compliance by design', as a way to either ensure automated enforcement of legal norms (including contractual obligations) or to ensure that legal protection (including contestation) is articulated into the data- and code-driven architectures of the onlife world. The first goes under the heading of 'legal by design', the second has been coined as 'legal protection by design'.

9

Legal Personhood for AI?

In 1942, science fiction author Asimov formulated his famous ‘Laws of Robotics’, in his short story *The Runaround* (included in his 1950 collection of short stories *I, Robot*):

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

These ‘laws’ raise more questions than they answer, which makes them a very interesting attempt to confront the unpredictability of autonomous computational systems.

The first type of question regards the sequence of the laws; as the cartoon in Figure 9.1 indicates, the sequence is not arbitrary.

The second type of question actually proves the point made by the cartoon; these laws (and the sequence of applying them) are not merely relevant for individual choice but implicate society as a whole. This also goes for the question of how society as a whole enables or restricts individual choice.

Asimov in point of fact articulated a fourth or zeroth law that was meant to precede the others:

A robot may not harm humanity, or, by inaction, allow humanity to come to harm.

By now the rise of autonomous systems (from connected cars and industrial robotics to search engines and fintech) has come to a point where the paradoxes that are implied in these laws become apparent. The Massachusetts Institute of Technology (MIT) has developed an online software tool that invited users to answer questions about the kind of choices that a fully

WHY ASIMOV PUT THE THREE LAWS OF ROBOTICS IN THE ORDER HE DID:

POSSIBLE ORDERING	CONSEQUENCES	
<ol style="list-style-type: none"> 1. (1) DON'T HARM HUMANS 2. (2) OBEY ORDERS 3. (3) PROTECT YOURSELF 	<p>[SEE ASIMOV'S STORIES]</p>	<p>BALANCED WORLD</p>
<ol style="list-style-type: none"> 1. (1) DON'T HARM HUMANS 2. (3) PROTECT YOURSELF 3. (2) OBEY ORDERS 	<p>EXPLORE MARS!</p>  <p>HAHA, NO. IT'S COLD AND I'D DIE.</p>	<p>FRUSTRATING WORLD</p>
<ol style="list-style-type: none"> 1. (2) OBEY ORDERS 2. (1) DON'T HARM HUMANS 3. (3) PROTECT YOURSELF 		<p>KILLBOT HELLSCAPE</p>
<ol style="list-style-type: none"> 1. (2) OBEY ORDERS 2. (3) PROTECT YOURSELF 3. (1) DON'T HARM HUMANS 		<p>KILLBOT HELLSCAPE</p>
<ol style="list-style-type: none"> 1. (3) PROTECT YOURSELF 2. (1) DON'T HARM HUMANS 3. (2) OBEY ORDERS 	 <p>I'LL MAKE CARS FOR YOU, BUT TRY TO UNPLUG ME AND I'LL VAPORIZE YOU.</p>	<p>TERRIFYING STANDOFF</p>
<ol style="list-style-type: none"> 1. (3) PROTECT YOURSELF 2. (2) OBEY ORDERS 3. (1) DON'T HARM HUMANS 		<p>KILLBOT HELLSCAPE</p>

Figure 9.1 By XKCD, This work is licensed under a Creative Commons Attribution-Non-Commercial 2.5 License

autonomous, self-driving car would have to make.¹ For instance, whether the car should prioritize its passengers when faced with the dilemma of either killing pedestrians or its passengers, or, whether it should decide such options by ranking people based on their age, their number, or other potentially relevant criteria. This raises yet other questions, such as whether such choices can be hardwired into the car's firmware by the manufacturer at its discretion or should be decided by the owner (which may be a car rental service) or the user (which may be anybody who actually 'drives' the car as a passenger). One could also imagine legislation where such choices are made by the legislature and imposed on developers, manufacturers, retailers, owners, and/or users.

¹ <http://moralmachine.mit.edu>.

A closer look at the reality of supposedly driverless cars, however, demonstrates two objections against the way the issue is framed. First, experts are not in agreement whether the level of autonomy that is assumed in the portrayal of these choices will ever be achieved. Some suggest that this type of robotics is running into a wall, due to the limitations of data-driven ‘intelligence’ in real-life scenarios and the risks its employment generates. In robotics, developers speak of ‘the envelop’ of a robot. ‘The envelop’ is usually designed simultaneously with the robot, to ensure its functionality and the safety of those interacting with it. Often this implies physically separate spaces for robots and humans, as the navigation of robots in a shared space generates substantial risk of harm. In point of fact, Rodney Brooks, a famous roboticist who designed an industrial robot that may be trusted in a shared space, predicts that self-driving cars will require separate lanes and roadblocks to reduce the risk of impact on human users of public roads.

Second, the issues are framed in somewhat naive utilitarian terms, defining the problem in terms of individual preferences that can then be aggregated and decided based on whatever the majority of a specific user-community prefers. Such a utilitarian calculus assumes that preferences are given, do not fluctuate over time, concern independent variables, and can be assessed out of context (based on a schematic depiction that restricts itself to specific details, abstracting from many other details). The framing plays around with the question of whether such choices are agent-dependent: will an agent’s choice about whose life must be prioritized depend on whether they are in control of the car’s behaviour, or on whether they may be the victim? Is the fact that the agent has a family relationship with the potential victim morally relevant, or should choices be made from behind a ‘veil of ignorance’ about such agent-dependent details? Is it a good idea to consider such questions to be a matter of individual preferences, similar to a taste for either red or white wine?

In the context of this book, the question we face is one of legal personhood rather than moral agency. In 2017, the European Parliament (EP) voted on a resolution, requiring the European Commission (EC) to address the potential of ‘civil law rules on robotics’.² The resolution was passed with 396 against thirteen votes, with eighty-five abstentions. Though the EC is not bound by

² European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)). At the end of this chapter we will return to the follow-up of the resolution, the Report on *Liability for Artificial Intelligence and Other Emerging Digital Technologies* (2019) from the Expert Group on Liability and New Technologies—New Technologies Formation, an independent expert group that was set up by the European Commission, <https://ec.europa.eu/transparency/regexpert/index.cfm?do=groupDetail.groupMeetingDoc&docid=36608>.

the resolution it must respond and explain if it does not act upon the recommendations it contains. Under point 31, the EP:

Calls on the Commission, when carrying out an impact assessment of its future legislative instrument, to explore the implications of all possible legal solutions, such as:

(...)

- f) creating a specific legal status for robots, so that
 - at least the most sophisticated autonomous robots
 - could be established as having the status of electronic persons
 - with specific rights and obligations,
 - including that of making good any damage they may cause, and
 - applying electronic personality to cases where robots make smart autonomous decisions or otherwise interact with third parties independently

This chapter will engage with the legal issues of autonomous systems, asking the question whether (and if so, under what conditions) such systems should be given legal personhood. Note that legal personhood can be attributed in the context of different legal domains: enabling legal personhood for corporations can, for example, be done for private law, while restricting criminal law liability to natural persons. It should be clear that criminal law liability with its emphasis on censure assumes a kind of moral agency that is not obvious in the case of current day autonomous systems. Strict liability in private law, however, would not necessarily be concerned with moral blame.

To investigate these issues, we will first discuss the concept of legal subjectivity and legal agency, followed by the concept of artificial agency, resulting in a first assessment of the potential of civil liability of autonomous systems.

9.1 Legal Subjectivity

In modern positive law, there are *two types of legal subjects*:

1. natural persons; and
2. legal persons.

Human beings are considered to be ‘natural persons’, though this should not be seen as something ‘natural’. In the past, human beings such as slaves and

women have been denied the status of legal subject, meaning they could not own property, not conclude contracts, they could not vote, or claim a right to privacy or freedom of expression. The decision that all human beings are legal subjects was a political decision that sprang from the idea that governments should treat each individual as deserving equal respect and concern (see above section 2.2 and 3.3).

Legal subjectivity is attributed by positive law, just like subjective rights (the rights of legal subjects) depend on objective law (the totality of rules and principles that decides what legal conditions result in which legal effect, see above section 3.1.3).

Apart from individual human beings, the law can and does attribute legal personhood to other entities, for instance to corporations, associations, and foundations or to municipalities and the state. So, both private and public bodies can qualify as legal persons if the legislature (or precedent in common law) attributes legal subjectivity to them. If so, they can act in law: own property, conclude contracts, and they can be held liable for damage caused under private law and they may even be charged with a criminal offence. However, whereas human beings are legal subjects under private, constitutional, and criminal law, this is not necessarily the case for legal persons such as a corporation. This will vary per jurisdiction; in some jurisdictions a corporation is a legal person under private law, but not under criminal law.

The concept of a person derives from the Latin *persona*, which means mask. A mask does two things: it enables one to play a role, and it shields the entity behind the mask. The mask thus provides its bearer with *positive freedom* (the role it can now play) and with *negative freedom* (warding off identification between the mask and its bearer). On the one hand, the ‘mask’ of the *legal persona* allows an entity to act in law (to create legal effect), and to be held liable; on the other hand, the ‘mask’ shields and thus protects that entity. The mask prevents identifying a person of flesh and blood with their role in law, thus ruling out that a person is defined by their legal status. This way the law leaves room for reinvention of the self. The idea of the *persona* is pivotal for the instrumental and protective role of law: it is an *instrument* where it enables an entity to act in law or to be held liable and it *protects* where it prevents equating legal status with the living person.

This raises the question of whether there are criteria that condition the attribution of legal personhood. Many authors believe that human beings are naturally legal subjects, whereas corporations are legal subjects due to a legal fiction. They are treated *as if* they are legal persons (as a *legal fiction*), whereas they are not ‘really’ persons or subjects. This has given rise to metaphysical musings about what distinguishes real from fictional persons. The problem with this perspective is that it overlooks the fact that the attribution of legal subjectivity always concerns an artificial construct. As John Dewey observed in a famous article on legal personhood, a legal fiction such as legal personhood is real even though it is artificial. Just like an artificial lake is a real lake, not an imaginary lake.

To emphasize that legal subjectivity is an artificial construct, based on a performative speech act that qualifies an entity as a legal subject, we should be reminded that at some point animals could be charged with a criminal offence; black people have been ‘regarded as beings of an inferior order’ with ‘no rights which the white man was bound to respect’ (*Dred Scott*);³ while, for instance, an unborn baby may be ‘regarded to have been born already as often as its interests require so’ (Article 1:2 Netherlands Civil Code).

The artificiality of legal personhood is related to the fact that legal subjectivity is by definition *attributed* by positive law (statute or common law) and cannot be assumed, while—in turn—the legal capacity of legal subjects can be restricted by positive law (for instance, in the case of minors, or in the case of guardianship).

Note that the terminology is such that the term *legal subject* is used for both *natural persons* and *legal persons*, whereas the term legal persons is only used for legal subjects that are not natural persons. As a consequence, legal persons always require *representation*; a corporation cannot act other than by way of its legal representatives. Clearly, if a legal person is liable under criminal law, it cannot be put in prison, though other punishments will apply (such as fines, closure of operations, or even termination of the organization).

All this should lead to the conclusion that, in principle, positive law can attribute legal personhood to whatever entity, depending on whether the legislature (or the common law) deems such an attribution necessary to protect legally relevant rights, freedoms, and interests.

³ *Dred Scott v. Sandford*, 60 U.S. (19 How.) 393 (1857).

9.2 Legal Agency

As to terminology, it makes sense to distinguish between:

1. a human person, used as a biological term (distinguishing humans from other animals, but also raising the question when a cyborg stops qualifying as a human person);
2. a moral person, used as a moral term (raising the issue of whether and if so, under what conditions, an artificial agent can be qualified as a moral person, capable of acting rightly or wrongly); and
3. a natural or a legal person, used as a legal term based on positive law (raising the question which animals or artificial agents would qualify for legal personhood, noting that this will involve a political decision).

These issues can also be framed in terms of agency instead of personhood, for example, in terms of moral agency, which is generally understood as the capability to engage in *intentional action*, which in turn assumes the capability of *giving reasons for one's actions*; or, in terms of legal agency, generally understood as the capability, attributed by law, to act in law and to be liable for one's own actions (legal subjectivity). Interestingly, however, there is a second meaning for the concept of legal agency, which refers to the capability, attributed by law, *to act on behalf of another* (acting as a proxy, a representative).

This second meaning of agency assumes a specific legal relationship between an *agent* and its *principal*, where the agent acts on behalf of a principal. This is usually based on a contractual relationship between the agent and its principal, on the one hand, and the agent and a third party, on the other hand, thus creating a contractual relationship between principal and third party, cf. Figure 9.2.




Figure 9.2 Agency relationship

For instance, a corporation that is running fashion shops in various locations may be represented by salespersons who actually sell clothing to visitors of the shop. In that case, the salesperson is the agent, the corporation is the principal. Note that under current law (in most jurisdictions) both the principal and the agent must be legal subjects for the third party to be bound by the actions of the agent. This already raises the interesting question of whether the corporation is bound if clothes are sold by an artificial agent (a bit of software that is part of a webshop that sells clothes online). The answer is yes, but this is based on the fact that such a software agent is considered a tool used by the corporation, not based on agency law.

Under agency law, it is crucial to establish *the authority of the agent* (that is, the extent to which the agent is allowed to act on behalf of the principal). We distinguish between the scope of the authority and its origin. As to the scope, the law differentiates between universal agents (that have authority for all acts), general agents (that have authority for all acts regarding a specific function), and special agents (with authority for one specific type of act).

As to the origin of the authority, the law differentiates between:

1. actual authority (implied or express);
2. ostensible, or apparent authority (estoppel); and
3. ratified authority (where the principal confirms authority despite the fact that the agent acted *ultra vires*, that is, beyond the stipulated authority).

An important question is whether the principal is liable for actions of an agent that acts *ultra vires*. In other words, does the legal effect of a contract with a third party, concluded by the agent on behalf of the principal, apply to the principal if the agent went beyond its authority and the principal did not ratify? The answer to this question depends on the following.

The principal is bound:

- if the third party was justified in trusting the agent to act within the scope of their authority, *and*
- the principal acted or omitted in a way that generated justified trust, *or*
- if the risk is for the principal, on the basis of generally accepted principles.

If these conditions do not apply the agent is liable.

9.3 Artificial Agents

Before moving deeper into the question of whether software or embedded systems can or should be qualified as legal persons, we need to define what is meant by an artificial agent.

Luc Steels (a renowned scientist working on AI), defines an agent as follows:

1. a system (a set of elements with relations amongst themselves and with the environment);
2. performing a function for another agent;
3. capable of maintaining itself.

He then differentiates between an automatic agent, that is self-steering on the basis of external laws, and an autonomous agent, that is both self-steering and self-governing.

In other work I have distinguished between automatic, autonomic and autonomous agency (where agency is defined as a combination of perception and the ability to act on what is perceived, while perception is informed by potential action):

1. automatic agency implies that the conduct of the agent is entirely predefined, for example, a thermostat or a smart contract;
2. autonomic agency implies that the agent is capable of self-management, self-repair, self-configuration, for example, a biological central nervous system, power management in a data centre, cooperating wireless sensor networks that 'run', for example, a smart home;
3. autonomous agency implies both consciousness and self-consciousness, meaning that the agent is capable of self-reflection, intentional action, argumentation, and the development of second order desires, for example, and notably human beings. Second order desires are desires about our desires, such as a desire not to desire smoking.

Steels' automatic agents would fit with my autonomic agency. Note that autonomic agency does not necessarily imply consciousness and many organisms, including conscious animals, would fall within its scope, whereas autonomous agency requires self-awareness in a way that escapes autonomic agency. It seems that moral personhood is contingent upon autonomous agency. If,

and to the extent that legal personhood would require self-consciousness, autonomic agents would not qualify. However, corporations that enjoy legal personhood are not self-conscious even if they may be represented by human beings that are.

This implies that there is no categorical legal answer to the question whether an autonomous computational system (usually an autonomic system in the above sense) should be given legal personhood.

That question is a political question that must be answered by a legislature weighing the advantages and disadvantages of such a move. There is, however, one caveat. Legal personhood that involves criminal law liability or constitutional rights, such as the right to privacy or non-discrimination, seem to require entities that can be called to account for their actions, which assumes a kind of self-consciousness. Interestingly, corporations can be made liable under criminal law and, for example, the ECtHR has found that corporations may have a right to privacy, despite corporations not having consciousness let alone self-consciousness. It is pivotal to acknowledge that legal personhood is always restricted, compared to the kind of full legal subjectivity enjoyed by natural persons, but it is also pivotal to recognize the fact that restricted forms of legal personhood have been attributed that seem to involve blame (criminal law liability) or the kind of freedom that is often at stake when human rights are violated (constitutional or fundamental rights).

The question to be answered when inquiring whether legal personhood should be attributed to artificial agents is a pragmatic question about:

1. what problem the introduction of such attribution solves;
2. what problem it doesn't solve; and
3. what problems it creates.

9.4 Private Law Liability

In this chapter, we will focus on the attribution of legal personhood to artificial agents that enables private law liability of such agents. If we follow the definition of Luc Steels, where an artificial agent acts on behalf of another agent, combined with the issue of an artificial agent acting on behalf of a natural or legal person, the following problem surfaces: under current law, to be

a legal agent implies being a legal subject, whereas an artificial agent would be a legal object, a tool, but not a legal subject. This means that an artificial agent cannot bind the legal subject on whose behalf it operates, other than as a tool. Many scholars have raised the question of an artificial agent that causes harm or damages in a way that was unforeseeable for its 'principal', as they fear that such unforeseeability will stand in the way of liability of the 'principal'.

In the case of machine-to-machine contracting with the help of software agents that are entirely determined by their algorithms, those employing the 'agents' can foresee what types of contracts will be concluded. In the case of machine-to-machine contracting with software agents that act autonomically (displaying, e.g. emergent behaviours), those employing them cannot foresee all the consequences. Insofar as this would imply that those employing such agents escape liability (as their own conduct may not have been wrongful, precisely because they could not have foreseen the harm), one could argue that victims would benefit if the agent itself could be held liable. To protect potential victims against suffering damages for which they cannot be compensated, artificial agents whose behaviour cannot be foreseen by those who employ them could be certified and registered as legal persons on the condition that they have access to funds that compensate potential victims in case of harm or damage. One could even imagine a prohibition of artificial agents with a propensity to cause harm or damage unless they are certified, registered, and either insured or provided with sufficient funds to compensate actual victims.

If the problem to be solved is that unforeseeable damage rules out liability of whoever employs the agent, we can foresee the following solutions:

1. *The agent can be seen as tool (as under current law):*
 - courts or legislatures could relax the requirement of intent or negligence on the side of whoever employs the tool (a move towards strict tort liability);
 - the law could deny validity to transactions that were generated by autonomic agents that are unpredictable (which might, however, stifle innovation).
2. *The artificial agent can be registered as a legal person (future law?):*
 - this would enable attribution of actual or ostensible authority to the agent, thus making its principal liable (raising the question of what's the difference compared to strict liability);
 - this would, however, also enable making the agents liable on their own account (certification, own funds, etc.) if, for instance, they overstep their authority.

The question of legal personhood for artificial agents clearly demonstrates that even if its attribution would solve some problems, it will create others. Many legal and other scholars warn that such attribution should not enable those who develop and employ artificial agents to outsource and escape responsibility, thus incentivizing them to take risks and externalize costs because they know they will not be liable.

In 2019, an Expert Group on Liability and New Technologies (set up by the European Commission), published its *Report on Liability for Artificial Intelligence*,⁴ in response to the resolution of the European Parliament, referred to in the introduction of this chapter. The Expert Group developed the following recommendations:

- A person operating a permissible technology that nevertheless carries an increased risk of harm to others, for example AI-driven robots in public spaces, should be subject to strict liability for damage resulting from its operation.
- In situations where a service provider ensuring the necessary technical framework has a higher degree of control than the owner or user of an actual product or service equipped with AI, this should be taken into account in determining who primarily operates the technology.
- A person using a technology that does not pose an increased risk of harm to others should still be required to abide by duties to properly select, operate, monitor, and maintain the technology in use and—failing that—should be liable for breach of such duties if at fault.
- A person using a technology which has a certain degree of autonomy should not be less accountable for ensuing harm than if said harm had been caused by a human auxiliary.
- Manufacturers of products or digital content incorporating emerging digital technology should be liable for damage caused by defects in their products, even if the defect was caused by changes made to the product under the producer's control after it had been placed on the market.
- For situations exposing third parties to an increased risk of harm, compulsory liability insurance could give victims better access to compensation and protect potential tortfeasors against the risk of liability.
- Where a particular technology increases the difficulties of proving the existence of an element of liability beyond what can be reasonably expected, victims should be entitled to facilitation of proof.
- Emerging digital technologies should come with logging features, where appropriate in the circumstances, and failure to log, or to provide reasonable access to logged data, should result in a reversal of the burden of proof in order not to be to the detriment of the victim.

⁴ See above footnote 2, at 3–4.

- The destruction of the victim's data should be regarded as damage, compensable under specific conditions.
- It is not necessary to give devices or autonomous systems a legal personality, as the harm these may cause can and should be attributable to existing persons or bodies.

It seems that the expert group seeks to solve problems caused by emergent behaviour and subsequent unpredictability of artificial agents by means of adaptation of the requirements of private law liability, without resorting to legal personhood for such agents. The main concern of the experts seems to be that those manufacturing, operating, using, or updating these agents must be held accountable for harm caused—to prevent hazardous employment of artificial agents. By ensuring that victims can hold to account those taking the risk of unpredictable artificial agents, this particular approach can stimulate innovation, as it will increase the reliability of artificial agents that are put on the market.

References

Re Asimov's laws of robotics

- Clarke, Roger. 1994. 'Asimov's Laws of Robotics: Implications for Information Technology'. *Computer* 27 (1): 57–66. <https://doi.org/10.1109/2.248881>.
- Pasquale, Frank A. 2017. 'Toward a Fourth Law of Robotics: Preserving Attribution, Responsibility, and Explainability in an Algorithmic Society'. *Ohio State Law Journal* 78 (5): 1243–1255.

Re legal subjectivity for non-humans

- Bryson, Joanna J., Mihailis E. Diamantis, and Thomas D. Grant. 2017. 'Of, for, and by the People: The Legal Lacuna of Synthetic Persons'. *Artificial Intelligence and Law* 25 (3): 273–91. <https://doi.org/10.1007/s10506-017-9214-9>.
- Chopra, Samir, and Laurens White. 2004. 'Artificial Agents: Personhood in Law and Philosophy'. In *Proceedings of the European Conference on Artificial Intelligence*, 635–39. IOS Press.
- French, Peter A. 1979. 'The Corporation as a Moral Person'. *American Philosophical Quarterly* 16 (3): 207–15.
- Hildebrandt, Mireille. 2011. 'Criminal Liability and "Smart" Environments'. In *Philosophical Foundations of Criminal Law*, edited by R.A. Duff and Stuart Green, 507–32. Oxford: Oxford University Press. <http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199559152.001.0001/acprof-9780199559152-chapter-22>.

- Koops, B.J., M. Hildebrandt, and David-Olivier Jacquet-Chiffelle. 2010. 'Bridging the Accountability Gap: Rights for New Entities in the Information Society?' *Minnesota Journal of Law Science & Technology* 11 (2): 497–561.
- Wells, Celia. 2001. *Corporations and Criminal Responsibility*. Vol. 2. [Oxford Monographs on Criminal Law and Justice]. Oxford and New York: Oxford University Press.

‘Legal by Design’ or ‘Legal Protection by Design’?

Policymakers, lawyers, and other folk often speak of ‘regulating technologies.’ This is an interesting phrase, because it can mean many things, depending on how you ‘read’ it. In the old days, most lawyers and policymakers would understand it in the sense of *technologies being the object of legal regulation*. The law can, for instance, impose requirements on the fabrication, design, sale, and use of cars, knives, guns, housing, office space, washing machines, toys, or medical instruments. These requirements may concern safety, privacy, or a technology’s potential to violate copyright, to disseminate child pornography, or to generate pollution of the environment. They may be aimed at protecting weaker parties, critical infrastructure, national or public security, or the environment. The default response that technologies are the *object* of regulation may, however, be changing.

The same phrase (‘regulating technologies’) can also refer to *technology as a ‘subject’ that is regulating human behaviour*, for example, by way of speed bumps, digital rights management (DRM) technologies, news feed algorithms that determine what news we perceive, and other default settings that determine our ‘choice architecture.’ Here, the object of regulation is not a technology but human behaviour. So, technology can be either the object or the subject of regulation (and maybe both), whereas law is usually only seen as a subject of regulation (that which regulates).

This may be about to change due to the pervasive effects of two types of technologies that impact the environment of the law: machine learning (ML) applications that, for example, decide a person’s credit worthiness or employability, and distributed ledger technologies (DLTs) that allegedly self-execute transactions and agreements without and beyond the law.

In this chapter, the focus will be on how ML and DLTs transform the environment of the law, the substance of legal goods (such as legal certainty, equality before the law, inalienability of personality rights, fairness, and human dignity) and the extent to which this affects legal protection.

One of the main challenges here concerns the regulatory effects of these novel technologies and the potential incompatibility of legal protection and *techno-regulation* (defined as the regulatory effects of a technology, whether or not intended).

10.1 Machine Learning (ML)

To understand the relevance of ML for legal protection, it may help to look at a very simply example, such as AB testing. Imagine that the provider of a website wants to 'optimize' it to achieve higher performance in terms of influencing its visitors' purchasing behaviours, their reading habits, or political.

To do so, the provider may employ software that enables the following process:

1. the current webpage is called version A;
2. its design is changed in a minimal way, for example, the colour or place of a button, the position of a text block, the type and number of clicks required to access other webpages within the site;
3. the slightly transformed page is called version B;
4. 50 per cent of visitors are directed to version A, the other 50 per cent to version B;
5. the software automatically measures the visitors' clickstream behaviours, possibly including those captured over the next day (possibly across various other websites based on tracking cookies);
6. the software calculates which version generated the more desirable behaviours;
7. the version that is more effective is then used as the default page;
8. the whole process is repeated with another slight change;
9. AB testing can be targeted at specific types of people or even be personalized.

Let's see if this qualifies as an example of ML. In his handbook on *Machine Learning*, Tom Mitchell recounts that:

A computer program is said to learn

- from experience E
- with respect to some class of tasks T
- performance measure P

if

- its performance at tasks in T ,
- as measured by P ,
- improves with experience E .

As to *type of task* T : this clearly sets out that machines do not learn anything if no task is defined. In this case, the task will be defined by the website ‘owner’, together with the software provider, because the definition of what counts as desirable behaviour needs to be translated into machine-readable language. A webshop may find increased purchasing behaviour desirable, though they may also formulate more complex tasks, based on a segmentation of the visitors: they may prefer to increase the purchasing behaviour of people who buy expensive products, or of people who are likely to buy more than one product over the course of a specified period of time.

As to *experience* E : note that the experience of this software is limited to clickstream behaviours of visitors of the page, even if they can be followed on other sites. It may be that their behaviours on other sites are not within the tracking-scope of the software provider (e.g. in offline shops or via another browser), whereas those unknown behaviours are actually more relevant for an inference about their preferences. The software’s experience, however, is necessarily limited to the available training data.

As to *performance metric* P : it may be that a simple performance metric, such ‘clicks on one product’, or ‘buys at least two products’, does not really say much about the preferences of the visitors, because these behaviours are instances of situated behaviour that depends on many other factors. These other factors may be more indicative of their preferences. To test both versions against each other, one may need to test six or seven different performance metrics to obtain a better picture of what qualifies as an accurate measure of achieving desirable behaviour.

10.1.1 Exploratory and confirmatory ML research design

AB testing can be done by way of an exploratory research design, meant to generate hypotheses about what kind of behaviour is more lucrative for the webshop. This implies recognition that such AB testing is a matter of real-time experimentation. As Hofman, Sharma, and Watts write:

254 'Legal by Design' or 'Legal Protection by Design'?

In exploratory analyses, researchers are free to study different tasks, fit multiple models, try various exclusion rules, and test on multiple performance metrics. When reporting their findings, however, they should transparently declare their full sequence of design choices to avoid creating a false impression of having confirmed a hypothesis rather than simply having generated one (3). Relatedly, they should report performance in terms of multiple metrics to avoid creating a false appearance of accuracy.

Claiming success based on such AB testing is a very bad idea, and usually amounts to what statisticians call p-hacking. For a reliable prediction one needs a confirmatory research design, that provides tested and testable hypotheses about the preferences of visitors. As Hofman, Sharma, and Watts write:

To qualify research as confirmatory, however, researchers should be required to preregister their research designs, including data preprocessing choices, model specifications, evaluation metrics, and out-of-sample predictions, in a public forum such as the Open Science Framework (<https://osf.io>).

As one can understand, providers of marketing software that enables micro-targeting or underlies behavioural advertising will not be inclined to deposit their research design, including pre-processing choices, at the OSF.

We may conclude from all this that:

- ML is used to influence or nudge people into behaviours that are desirable from the perspective of whoever pays for the software; and
- such software may not be as effective as some may either hope or fear.

10.1.2 Implications of micro-targeting

Instead, the result of micro-targeting based on flawed research design may be that visitors of websites are confronted with a personalized choice architecture that is meant to lure them into what others find desirable behaviour.

This has two unintended consequences:

1. a fragmented public space that might algorithmically favour extreme content to hold onto people's attention; and

2. undesirable discrimination based on data points that systematically disadvantage certain categories of people.

These consequences are not necessarily envisaged by developers or users of the software; they are brought about by mistaking—potentially crappy—exploratory research design for robust confirmatory research design.

This raises issues for legal protection. For instance, the mining and inferencing of behavioural data may interfere with *specified fundamental rights, such as privacy, data protection, non-discrimination and freedom of expression*. Behavioural data are often personal data and the mining of such data may infringe the privacy of those unaware of the rich profiles that can be built from such data, often combined with features that are inferred from such data. This may be in direct violation of the fundamental right to data protection, depending on how the data is mined and shared, on what ground, and with what purpose (see above, section 5.5.2). Based on micro-targeting, the mining and inferencing of behavioural data may also violate the freedom of expression, since this right includes the freedom to receive information free of censure. Micro-targeting based on AB testing could shield information from certain people, because there is no added value for the website owner in providing them with such information. We have entered the era of ad-driven-content, where the algorithms that infer what content is most conducive to attracting visitors may be prioritized in order to increase ad revenue. The use of ‘low hanging fruit’ to train ML algorithms will easily result in all kinds of unwarranted bias, due to the bias that is inherent in the so-called ‘training data’. Even if the right kind of data is available, the choice of the feature space, the hypothesis space, the task that is formulated, and the performance metric that is chosen may result in a biased outcome that systematically discriminates against people based on their race, ethnicity, religion, political preferences, gender, or sexual orientation.

An example of such bias is the proprietary COMPAS software, sold by Equivant (formerly Northpointe), where COMPAS stands for Correctional Offender Management Profiling for Alternative Sanctions. COMPAS is used by courts in the United States to assess the risk that an offender will recidivize (i.e. commit another offence after being released). This risk co-determines the parole or sentencing decisions. The risk score is based on a limited number of data points that have been found to correlate with re-offending. COMPAS is the result of an ML research design that tested 137 features to infer which six features were actually predictive. After Julie Angwin conducted own research

on similar training data, she claimed that COMPAS discriminates against people based on their race.

More precisely, Julie Angwin found that:

- within the set of offenders that did not recidivize, the error rate for black persons may have been as high as that for white persons, but the error for black persons meant they were wrongly given a higher risk score;
- whereas the error for white persons meant they were wrongly given a lower risk score.

According to Equivant, this was the result of the fact that black persons on average had a higher risk of reoffending. Equivant accused Angwin of methodologically flawed methodologies, implying the laws of statistics were responsible for the disparate outcome of the risk score. As a use case, the accusation of racial discrimination has generated a flood of scientific literature on fairness in ML, underpinning requests for transparency and accountability, basically demanding that business and government employs FAT ML (fair, accountable, transparent machine learning applications).

The literature demonstrates that many different definitions of what qualifies as fair ML are possible, leading to different research designs. For instance, in the case of COMPAS, one could argue that fairness requires that the 'learner' is trained to come up with a risk score that does not result in disparate errors for black and white persons who do not recidivize. The COMPAS case returns in more detail in section 11.3.2.1.

10.1.3 Implications of micro-targeting for the rule of law

The second issue for legal protection concerns the extent to which decisions based on ML-inferences violate *core principles of the Rule of Law*, such as transparency and accountability.

Or, more precisely:

1. the explainability of the decision-making process;
2. the justification of the decision; and
3. the contestability of the decision.

The second and third requirements concern the decision. In public administration, decisions must be taken in accordance with the legality principle, meaning that the justification must be based on law and citizens have a right to contest the decision in a court of law (see above, section 3.1.2). In the private sector, however, the freedom of contract and the freedom to dispose of one's property as one wishes may provide the justification. These freedoms, however, are restricted, for instance due to the prohibition to discriminate in the context of employment, or to discriminate based on gender or race. Both in public administration and commercial enterprise, ML-based decision-making may incur *invisible discrimination* that is actually prohibited, for instance based on race. Such discrimination will often be unintended and invisible because it is based on a concerted set of features that correlate with race and therefore act as proxies for race. This means that such discrimination need not be based on a deliberate attempt to use race as a relevant feature; even if one removes race as a feature altogether, the proxies will probably sustain the discrimination.

Legally speaking, this may be qualified as *indirect (or disparate) discrimination*, which is often explicitly defined and prohibited by law (unless justified). What matters here is that without explainable ML, it may be very difficult to check the extent to which discrimination occurs.

Apart from prohibited discrimination, decision-making based on applied ML may have other repercussions.

Imagine that the risk profile that is applied to a person is based on:

- the average risk in a specified class of people,
- whereas that average risk does not apply to each member of that class.

In that case, individuals are basically treated on the basis of a score that probably does not apply to them. Even if such classification of individuals does not involve prohibited discrimination, it may be seen as unfair. For instance, on average women may have a risk of one out of eight to suffer from breast cancer. Depending on a woman's age, the occurrence of breast cancer in her ancestry and family, her lifestyle, and other factors, her risk will stray from 'one out of eight', to a potentially much higher or lower risk. Treating each and every woman as if her risk is one out of eight would therefore be unwise, and in the case of, for example, a health insurance premium

one might argue this is unfair. This explains why the explainability of decisions based on the application of ML has become a serious issue of legal protection.

In terms of the General Data Protection Regulation (GDPR), personalized targeting based on ML would most often fall within the scope of Article 4(4):

'profiling' means any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements;

Article 21 GDPR stipulates that data subjects have a 'right to object' to profiling that is based on the grounds of Article 6(e) and (f), that is, based on a public task or public authority, or on the legitimate interest of the data controller, as well as a right to object to profiling 'to the extent that it is related to direct marketing'.

Next to this, data subjects have a 'right not to be subject' to profiling when this results in automated decision-making that significantly affects data subjects (Article 22). In section 10.3.3.3, we will explain to what extent the right not to be subject to automated decisions provides 'legal protection by design' against biased ML applications. Note that this right is not only applicable to profiling but also to other types of automated decisions, such as those involving self-executing code.

10.2 Distributed Ledger Technologies (DLTs), Smart Contracts, and Smart Regulation

As the development and usage of DLTs and/or Blockchains are in full flux, so is the terminology.

Sidestepping discussions on the correctness of either term, we will use the term DLT to cover the whole range of technologies framed as:

1. distributed databases (ledgers)
2. that store transactions based on

3. decentralized infrastructures (the core code)
4. that enable self-executing code, based on
5. a specific combination of security technologies (notably hashing and encryption)
6. that incentivize ‘miners’ or ‘validators’ to partake in a reasonably trustworthy consensus mechanism
7. that supposedly ensures the integrity of the data stored in the ledger, and of the sequence of such storage.

DLTs are often promoted as providing ‘*trustless*’ computing that enables immutable, transparent, and secure storage of transactions, with a guarantee against *ex-post* manipulation of previous transactions, thus ensuring the integrity of both the sequence and the content of the transactions (where the integrity of the sequence protects against ‘double spending’). Often DLTs are ‘sold’ as enabling *disintermediation*, meaning that users need not connect with a traditional institution (such as banks) to engage in trustworthy transactions with parties they do not know or do not trust. The idea is that the ledger allows them to interact with others in a fully transparent way, with certainty that neither the other party nor any third party can manipulate stored transactions. In a sense, the promise is that the technology can take over the role of a trusted intermediary by way of a fully predictable sequence of events that self-executes tamper-free transactions.

Before unpacking these claims, it is crucial to distinguish between *public and private* and between *non-permissioned and permissioned DLTs*, as well as their combinations.

The difference between public and private DLTs can be defined as depending on who can ‘read’ the content, and the difference between permissioned and non-permissioned can be defined as depending who can *add* or ‘write’ new content. Bitcoin builds on public non-permissioned DLTs, meaning that anybody can check the content and submit new content. By now, commercial enterprises, financial institutions, as well as government agencies probe business cases for DLTs, often resorting to private permissioned versions that lack part of the lure of a decentralized system, because with private permissioned DLTs only a specified set of players is allowed to read and write on the ledger.

Let us note up front that this means that private non-permissioned DLTs basically require users to trust:

1. the traditional intermediary that employs the DLT; and
2. those who write the code for a particular type of transaction; and
3. those who write the protocols that constitute the infrastructure ('core developers').

Taking into account that most users do not understand computer code, such DLTs basically reinforce the role of the institutions that employ them; they require more trust, not less, and they certainly do not achieve disintermediation.

10.2.1 Smart contracts and smart regulation

For this chapter, the relevance of DLTs concerns so-called *smart contracts* and *smart regulation*, that is, the use of DLT to self-execute either an agreed contract or a specified policy based on regulatory competence. As to the first, we can think of a contract of sale that self-executes once triggered (when the system detects payment it transfers the object, or the other way around). Note that this may work perfectly if both the payment (e.g. cryptocurrency) and the assets (e.g. an electronic proof of ownership) are within the system (often referred to as being on-chain). Off-chain payments or off-chain transfer of assets, however, will require the use of 'oracles', that is, software applications that interface between the ledger and the real world, or other systems.

As to the self-execution of regulatory policies this assumes that a competent authority *translates* its policy into machine readable code (an act of interpretation) and *defines* what kind of data-input triggers the execution of the code (another act of interpretation).

Some have observed that this conflates legislation with its execution and even with adjudication (in case of disagreement about the content of the contract). This would mean that the checks and balances of the rule of law, notably the separation of the powers of legislation, administration, and adjudication, are disrupted. This, in turn, would require new types of safeguards (legal remedies) to enable the contestation of the ensuing decisions—thus ensuring that smart regulation and smart contracts remain 'under the rule of law'.

A quick round-up of critique regarding some of the claims made about DLTs, notably with regard to smart contracting and smart regulation:

Immutability:

1. if parties agree on the code (and if the code is not corrupted or otherwise disabled), their agreement will be executed without recourse to remedies or re-interpretation. One could argue that the immutability of self-executing code entails legal certainty, though changing circumstances may result in the opposite, precisely because the code is not adaptive;
2. if a party does not understand code and agrees to oral or written communication that differs from the code, the immutability becomes a problem and will certainly not deliver legal certainty;
3. in the case of a permissioned DLT the immutability may be overruled, depending on the governance structure (as this may be distributed but will not be decentralized).¹

Trustless computing:

1. if parties do not know each other but wish to engage in transactions, a smart contract is often said to enable trustless computing, to the extent that the protocols of the platform and the code of the contract are trustworthy and do what parties legitimately expect;
2. parties are basically asked to trust that the protocols of the underlying infrastructure are trustworthy, and the program language aligns with the intent expressed;
3. in the case of permissioned private DLTs, users are required to trust (1) those who control the DLT, (2) the protocols that form its 'constitution', and (3) the code that is run on their behalf or on behalf of the other party.

Transparent transactions:

1. to the extent that parties have access to the source code of the infrastructure (public DLTs) and to the programming code (the smart contract itself), and to the extent they can understand the code, there is transparency;
2. if parties have no access to the code (private DLTs), or do not understand code, transparency cannot be assumed.

¹ Here, I use 'distribution' to refer to the physical location of the code or the data, and 'centralization' to refer to the power structure (who decides what).

Security:

1. if all works as hoped for, the execution of the contract is secured;
2. if the protocols and/or code are sloppy, if new bugs appear, in the case of a so-called 51 per cent attack, or if the miners/validators stop maintaining the system, the contract and/or the whole system may be hacked and/or dissolve.

Anonymity:

1. depending on how parties access the smart contract ecosystem, they may remain anonymous or at least pseudonymous;
2. transparency in public DLTs may imply that anonymity is an illusion, also considering the use of, for example, behavioural analytics to re-identify users.

Safety:

1. to the extent that the underlying system and the smart contract itself operate as agreed, the transaction could be called safe;
2. if circumstances change, requiring adaptation of the contract or decision, the self-executing nature of the code may create unsafe outcomes for users, especially if they cannot identify or sue whoever is liable (as the provider of the DLT, the contracting party, or the government agency may, for instance, be in another jurisdiction);
3. if either the underlying system or the smart contract code is hacked, if off-chain input is incorrect, or if the provider cannot be held liable, one or more parties to the contract may lose their input.

Correctness:

1. to the extent that off-chain input is correct, the on-chain execution of the contract will be executed correctly (as long as the code does what the parties agreed to);
2. to the extent that off-chain input is incorrect, the error or false input is automated (and, due to the immutability, this may be hard to correct).

From the perspective of law, the employment of DLTs raises many questions. In the context of this chapter, I focus on whether operating self-executing code via a DLT must be seen as 'legal by design' or as 'legal protection by design' (preparing the ground for the topic of section 10.3).

Do smart contracts or smart regulations guarantee that the behaviour of parties to the contract or of addressees of regulation is ‘legal by design’ or ‘legally compliant by design’? To prepare the ground, I will first discuss the question whether smart contracts are *contracts in the legal sense* (section 10.2.2), and whether smart regulation is *law in the legal sense* (section 10.2.3).

10.2.2 The legal status of ‘smart contracts’ under private law

As to contracts in the legal sense, we need to investigate what legal conditions must be fulfilled for ‘something’ to qualify as a legally binding contract. These legal conditions can be found in private law, which—in Europe—is mostly national law, as there is no binding European private law. I refer to section 3.2.2, where some of the basics of a valid contract were discussed, based on Dutch private law.

Though other jurisdictions may have different legal conditions some of the underlying assumptions remain the same.

- First of all, an *obligatory agreement* is a more-sided act where parties aim to establish specified legal effects, such as a legal obligation to pay a price in exchange for the transfer of property or the provision of a service.
- In the common law, a contract requires ‘consideration’ (tit for tat) to be valid.
- The intent to be bound by the contract can be *inferred* from the declarations of the parties, though sometimes it can also be inferred from their actions—if such actions have generated the legitimate expectation that one has consented to the contract.
- In most—if not all—jurisdictions, a valid contract requires a *sufficiently specified offer* by one party that is *accepted* by the other party.
- If the acceptance was mistakenly inferred from certain behaviours, whereas in fact there was no acceptance, the contract would be considered *void* (as one of the constitutive conditions does not apply).
- If the offering party, however, legitimately inferred acceptance from the other party’s behaviour, the contract may nevertheless be *valid*.

Most jurisdictions have safeguards in place in case acceptance is based on duress or undue influence, mutual mistake, or fraud. If this can be proven, the contract

becomes *voidable*, depending on the request of the party that wishes to 'undo' the contract.

In most jurisdictions, there are no formal requirements for contracts in general, which means they can be concluded in whatever way (speech, writing, shaking hands, real-time exchange of a good and the payment).

Specific contracts, such as the sale of real-estate, do have formal requirements (e.g. of a deed) which usually involves a trusted third party (e.g. a notary public).

Does a smart contract qualify as a legal contract?

Based on the above, there are at least three issues:

1. Can we assume that sending a message to a smart contract (code on the ledger) implies the will to be bound (and thus acceptance of an offer)?
2. Does computer code count as an expression of the content of a contract (and thus as a sufficiently specified offer)?
3. Can a party invoke voidability because they cannot read the code?

The fact that most contracts have no formal requirements could be used as an argument that sending a specific message to the code on the DLT may count as an expression of one's intent to enter into the contract as defined in the code. However, the jury is still out on whether computer code counts as an expression of the content of a contract just like a written contract supposedly does. To count as such an expression, the code must be sufficiently determinate for both parties to understand the legal effect of the contract (i.e. the legal obligations it generates).

If the accepting party does not read code, they can either:

- argue that they did not accept the content of the code because their legitimate expectations about that content—as inferred from negotiations, advertising, or other expressions by the offering party—do not match the code, which means the contract is void; or
- argue that the contract is voidable because of, for example, mistake or fraud.

If we assume that the contract is valid, we still need to look into the legal effect of a valid contract, because *in most jurisdictions such legal effect is not limited to the literal wording of the contract.*

It often extends to:

- what both parties should reasonably expect, considering the circumstances,
- while a number of legal constraints may apply that co-determine the content of the contract.

The latter constraints may derive from either private or public *mandatory law* (see section 3.1.2 and 8.1.1), which cannot be overruled by contractual stipulations (whether in speech, writing, or code). To build flexibility into a contract or a policy, they often contain concepts with an *open texture* that leave parties or competent authorities some room to adapt the contract to concrete circumstances that cannot all be foreseen. Think of terms such as reasonably, timely, state of the art, or trustworthy, which can only be interpreted in the light of the circumstances that parties confront when performing the contract. Unforeseen changes in circumstances may have an impact on the content of the ensuing legal obligations, as when one party can claim *force majeure*. Whereas the ‘smart contract’ will self-execute, *force majeure* may overrule the obligation to perform the contract, meaning the execution may have to be undone (which may be impossible and/or the party that benefits may not be identifiable, or in a far-out jurisdiction, meaning they cannot be sued).

All this also happens to ‘normal’ contracts, and with ‘normal’ decision-making in public administration, but it is crucial to highlight that smart contracts and algorithmic decision-making in the sense of smart regulation do not necessarily solve these problems and may indeed create extra problems, precisely due to *the non-adaptive nature of self-executing code*. Those who wish to remedy these new problems by creating adaptive code must realize that this implies foreseeing all possible future scenarios, which is by definition not possible. Though the attempt to foresee changing circumstances may prevent some problems, it still implies that legislation (a contract can be seen as legislating how parties should act), execution (a contract should clarify what counts as a performance), and interpretation (the meaning of a contract depends on the circumstances) are *all predetermined upfront by whoever writes the code*. This somehow scales the past while it freezes the future.

Legal scholar Allen argues that smart contracts will be part of what he has called the ‘contract-stack’, which involves speech acts, behaviour, written documents, deeds, electronically signed documents, and—potentially—also

self-executing code. This implies that contract law will be transformed to accommodate the use of self-executing code, for example, by way of legislation, case law, and doctrinal innovation. Similar arguments can be made for smart regulation, which could similarly be seen as a 'regulatory-stack', involving legislative Acts that grant regulatory competences, policy documents, government agency's behaviour patterns, decision-making processes and procedures, and—potentially—also self-executing code.

10.2.3 The legal status of 'smart regulation' under public law

With the term 'regulation' I refer to rules promulgated by public administration, or by independent supervisors that have been instituted by an Act of the legislature (usually called 'regulators' in the United States and the United Kingdom, e.g. the Federal Trade Commission; in the EU we can think of the EDPS or the national DPAs).

Such rules are either:

- part of an explicitly attributed *competence to create and impose rules*; or
- a way to provide transparency about *how a regulator will make use of its discretionary competence* (in that case, those rules form a policy).

Many government decisions affect individual citizens, such as the granting of a permit, social security, or a decision on taxation. Many of the arguments provided in the previous section can be repeated here, and do not merely apply to implementation via DLTs but also to other forms of algorithmic (automated) decision-making. It simply means that the relevant rules are interpreted and translated into non-ambiguous code, to enable their self-execution.

As with private law contracts, smart regulation will necessarily be overinclusive and underinclusive (or both), due to its lack of adaptive flexibility.

The need to formalize will—in a sense—freeze future responses into a template that necessarily overlooks changing circumstances and may not reflect developments in case law, which could result in the code violating rights instead of enforcing compliance. In that respect, it is crucial to remember that

these rules and policies, as well as their machinic automation, fall under the rule of law.

Instead of understanding ‘smart regulation’ as a kind of law, it is therefore better understood as public administration.

This means that these rules and policies, as well as their machinic translations, must at some point be contestable in a court of law. Those subject to decisions based on smart regulation should be capable of requesting a justification of the decision in accordance with the legality principle. Note, however, that a *justification* is not equivalent to an *explanation*, which rather serves as a means to make the decision contestable as to its justification.

10.3 ‘Legal by Design’ or ‘Legal Protection by Design?’

Some authors claim that self-executing code could be used to ensure that the conduct of legal subjects will be ‘legal by design’ (LbD). What they mean to say is that one can interpret the content of a contract, the content of policy guidelines, or even the content of legislation such that it becomes amenable to a translation into computer code. So-called ‘Turing complete languages’ have been developed in the realm of DLTs, to write ‘smart contracts’ that—as we have seen in section 10.2—supposedly self-execute whatever has been agreed by the parties. One can imagine similar attempts to ensure compliance at the level of regulatory rules.

10.3.1 Legal by design (LbD)

LbD is a subset of what other authors have termed ‘techno-regulation’. This refers to the fact that technologies often induce or inhibit and enforce or preclude certain types of behaviours, which has a *de facto regulatory effect*.

As mentioned in the introduction to this chapter, such regulatory effects can be:

- the result of *deliberate design* of a technology (requirements that specify which functions must be engineered); or

- the *unintended result of design choices* made with other intentions, or of *unforeseen usage* of the technology.

In the latter cases, we speak of side-effects, though we should take note that such side-effects may be more prominent or influential than the intended effects.

LbD is a *specific subset of techno-regulation* that is:

1. the result of deliberate design choices, where
2. those choices aim to ensure compliance with legal obligations by way of technical enforcement.

LbD involves two steps:

1. it involves a specific (non-ambiguous) interpretation of the relevant legal norm; and
2. it involves the translation of that interpretation into a programming language.

Note that these steps can be analytically distinguished, but may be conflated in practice (thus hiding the act of interpretation). Due to the need to select an interpretation that can be translated into unambiguous machine language, such interpretations may be *overinclusive* or *underinclusive* compared to the relevant legal norm.

For example, a legal obligation for an employee to drive a truck from A to B within a reasonable time scale could be part of a smart contract between an employer and an employee. As the performance of the contract takes place off-chain, an oracle must be put in place to provide clear signals about whether or not this legal obligation has been fulfilled. To define what performance counts as 'reasonable', taking into account various types of circumstances, the contract must be interpreted beforehand and translated into a set of input variables for the oracle. As discussed in section 10.2.2, 'reasonableness' is not a subjective concept under contract law as it will have to be interpreted in line with relevant case law, while taking account of the unique circumstances of the case at hand. This makes it highly unlikely that a smart contract can be equated with 'legal compliance by design', due to the rigidity of the behaviour of computer code compared to the adaptiveness of the meaning of natural language.

Another example could be that the legally allowed level of pollution caused by a car is integrated into smart regulation that rules out delivery of non-compliant cars by the car manufacturer. To enable this, however, the cars must be tested before leaving the factory, which necessarily disregards the actual pollution caused on the motorway. This, again, implies that there is no absolute guarantee that the car manufacturer is 'legally compliant by design'.

In both examples, LbD seems to be *an inept term* for what is actually achieved. As long as this is kept in mind, incorporating checks and balances (including legal remedies if the lawfulness is contested), smart contracts and smart regulation may nevertheless contribute to (though not guarantee) compliance.

10.3.2 Legal protection by design (LPbD)

Legal protection by design (LPbD) is another matter. It does not aim to guarantee enforcement of whatever legal norm, but rather aims to ensure that legal protection is not ruled out by the affordances of the technological environment that determines whether or not we enjoy the substance of fundamental rights.

The term 'legal', here, involves two important requirements of law in the context of a constitutional democracy:

- the scope of LPbD should be determined by way of *democratic participation*, for instance in the context of participatory technology assessment and involvement of the democratic legislature;
- those subject to such LPbD should be able to *contest its application in a court of law*.

Techno-regulation in general does not include these requirements and neither does LbD, which is often focused on excluding the involvement of trusted third parties. These two requirements thus distinguish LPbD from other types of 'by design' solutions, for instance 'value sensitive design' or 'privacy by design'. The latter are often proposed as *ethical requirements*, which is problematic for two reasons. First, as ethical norms cannot level the playing field, companies that apply such ethical design may be pushed out of the market. Second, ethical 'by design' approaches make protection dependent on the ethical inclinations of those who develop and market the choice architecture of

citizens, instead of demanding that such choice architecture must meet minimum standards that provide effective and practical protection. For readers interested in the confrontation of law and ethics, see Chapter 11.

10.3.3 LPbD in the GDPR

10.3.3.1 Data protection impact assessment

Three interesting examples of LPbD can be found in the GDPR. First, the legal obligation to conduct a data protection impact assessment (DPIA) in Article 35, which is compulsory if the introduction of a new technology is likely to present a high risk to the rights and freedoms of data subjects:

1. Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data. A single assessment may address a set of similar processing operations that present similar high risks.

(...)

3. A data protection impact assessment referred to in paragraph 1 shall in particular be required in the case of:
 - a) a systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person;
 - b) processing on a large scale of special categories of data referred to in Article 9(1), or of personal data relating to criminal convictions and offences referred to in Article 10; or
 - c) a systematic monitoring of a publicly accessible area on a large scale.

(...)

7. The assessment shall contain at least:
 - a) a systematic description of the envisaged processing operations and the purposes of the processing, including, where applicable, the legitimate interest pursued by the controller;
 - b) an assessment of the necessity and proportionality of the processing operations in relation to the purposes;
 - c) an assessment of the risks to the rights and freedoms of data subjects referred to in paragraph 1; and
 - d) the measures envisaged to address the risks, including safeguards, security measures and mechanisms to ensure the protection of personal data and to

demonstrate compliance with this Regulation taking into account the rights and legitimate interests of data subjects and other persons concerned.

(...)

11. Where necessary, the controller shall carry out a review to assess if processing is performed in accordance with the data protection impact assessment at least when there is a change of the risk represented by processing operations.

Recital (75) adds some considerations concerning the question what constitutes the *likelihood of a high risk to the rights and freedoms of natural persons*.

The risk to the rights and freedoms of natural persons, of varying likelihood and severity, may result from personal data processing which could lead to physical, material or non-material damage, in particular:

- where the processing may give rise to discrimination, identity theft or fraud, financial loss, damage to the reputation, loss of confidentiality of personal data protected by professional secrecy, unauthorised reversal of pseudonymisation, or any other significant economic or social disadvantage;
- where data subjects might be deprived of their rights and freedoms or prevented from exercising control over their personal data;
- where personal data are processed which reveal racial or ethnic origin, political opinions, religion or philosophical beliefs, trade union membership, and the processing of genetic data, data concerning health or data concerning sex life or criminal convictions and offences or related security measures;
- where personal aspects are evaluated, in particular analysing or predicting aspects concerning performance at work, economic situation, health, personal preferences or interests, reliability or behaviour, location or movements, in order to create or use personal profiles;
- where personal data of vulnerable natural persons, in particular of children, are processed; or
- where processing involves a large amount of personal data and affects a large number of data subjects.

Article 35 basically requires controllers to err on the side of caution by foreseeing risks to the rights and freedoms of natural persons. One could qualify this as the introduction of *the principle of precaution* in data protection law. Note that the assessment does not merely regard potential violations of the rights and obligations stipulated in the GDPR but focuses on 'rights and freedoms' in a more general sense, which links up with the goal of the GDPR as formulated in Article 2.2: '[t]his Regulation protects fundamental rights and freedoms of natural persons and in particular their right to the protection of personal data'. Moreover, the assessment of such a risk is not limited to data subjects but refers to 'natural persons', which includes individuals that run a

risk of being discriminated against even though their personal data are not (yet) being processed.

10.3.3.2 Data protection by default and by design (DPbDD)

Article 35.7(d) clearly indicates that a DPIA incorporates an assessment of the need for data protection by default and by design (DPbDD), as it requires an inventory of 'the measures envisaged to address the risks, including safeguards, security measures and mechanisms to ensure the protection of personal data and to demonstrate compliance with this Regulation taking into account the rights and legitimate interests of data subjects and other persons concerned'. This brings us to Article 25, which requires to design systems that process personal data in such a way that data minimization is achieved by default, while incorporating all other GDPR obligations into the design of the system:

1. Taking into account the state of the art, the cost of implementation and the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for rights and freedoms of natural persons posed by the processing, the controller shall,
 - both at the time of the determination of the means for processing and at the time of the processing itself,
 - implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation,
 - in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects.
2. The controller shall implement appropriate technical and organisational measures for ensuring that,
 - by default,
 - only personal data which are necessary for each specific purpose of the processing are processed.

That obligation applies to

- the amount of personal data collected,
- the extent of their processing,
- the period of their storage and
- their accessibility.

In particular, such measures shall ensure that by default personal data are not made accessible without the individual's intervention to an indefinite number of natural persons.

Here again, we can observe a requirement to err on the side of caution, basically echoing longstanding security principles, such as ‘select before you collect’. In paragraph 2, for instance, we read that *technical and organizational measures* must be in place to ensure that only data that is necessary for each specific processing purpose is processed (data minimization and purpose limitation). Though ‘privacy by design’ has deep roots in privacy engineering communities, the big difference with the new legal obligation is that this is no longer a matter of the arbitrary preferences of a company or public body that is ‘being ethical’ about their processing operations.

Though DPbDD is not to be taken lightly, it does not require what is not feasible. The obligation takes into account ‘the state of the art, the cost of implementation and the nature, scope, context and purposes of processing’ (first paragraph), meaning that measures must be doable, also in light of the business model. However, this does not mean that anything goes if the business model does not fly without taking *disproportionate risks* with the rights and freedoms of natural persons. Here again, as with the DPIA, those risks must be taken into account when designing (engineering) the processing operations. The proportionality depends on ‘the risks of varying likelihood and severity’, meaning that the higher the risks the more protection must be implemented ‘by design’.

Clearly, both the DPIA and DPbDD take a so-called ‘risk approach’ to the protection of personal data. Though some have interpreted this as a sign that the EU legislature favours a cybernetic understanding of risk and regulation over a rights-based approach, it seems more likely that the risk approach aims to introduce some lawfully required *precaution* on the side of data controllers, to sustain and enable an effective and practical protection of the rights and freedoms of natural persons.

When reading the carefully crafted, balanced, and reasonably complex requirements to embed relevant legal norms in the architecture of personal data processing, it is evident that neither the DPIA nor DPbDD aim to produce processing systems that are ‘legal by design’. Instead, they warrant and introduce legal obligations to embody ‘legal protection by design’ in technical systems that would otherwise render the protection of an individual’s rights and freedom illusory.

10.3.3.3 Automated decisions

This brings us to a *third example of LPbD* in the context of the GDPR that is highly relevant for both ML applications and DLTs, as it targets the implications of automated decisions. Article 22 GDPR reads:

The data subject shall have the right not to be subject to a decision

- based solely
- on automated processing, including profiling,
- which produces legal effects concerning him or her or
- similarly significantly affects him or her.

The legal effect of the four legal conditions (two of which are alternative), is a prohibition. Even though this prohibition is formulated in a rather complicated way, the European Data Protection Board (EDPB, formerly Article 29 Working Party) has clarified that this 'right not to be subject to' must be understood as a prohibition.² Note that each term in this set of legal conditions requires an act of interpretation that is not obvious in the light of technologies such as ML and DLT. For instance, which of the decisions taken by machines in the course of a machine learning operation qualify as a decision in the sense of Article 22.1: the decision of an algorithm to adept weights within a neural net, where such a decision will result in a refusal to provide credit? or, the decision to select four of the nineteen features that have some impact on a specified health risk, where such a decision results, for example, in a person being advised to undergo an operation or in a person being charged with tax fraud? Does 'solely' refer to machine decisions that directly affect a data subject (e.g. online acceptance of health insurance), or also to decisions that have been prepared by a software program but are 'stamped' by a human person who, however, does not understand how the system came to its conclusion and cannot explain to the data subject why she was not, for example, selected for a job interview? The EDPB finds that '[t]he controller cannot avoid the Article 22 provisions by fabricating human involvement'.³ Does the fact that automated processing is qualified as 'including profiling' imply that 'smart contracts' that do not involve profiling in the sense of Article 4(4) do not fall within the scope of Article 22? Note that English grammar answers that question, due to the fact that a comma is inserted after processing (check the rules for restrictive and non-restrictive modifiers).

When does a decision produce legal effect? The EDPB clarifies that this is the case if the decision 'affects someone's legal rights, such as the freedom to associate with others, vote in an election, or take legal action. (. . .) affects

² Article 29 Working Party WP251rev.01, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, at 19.

³ Ibid. at 21.

a person's legal status or their rights under a contract'.⁴ Any other 'similarly significant effect' also results in a prohibition, for example, as the EDPB writes:⁵

For data processing to significantly affect someone the effects of the processing must be sufficiently great or important to be worthy of attention. In other words, the decision must have the potential to:

- significantly affect the circumstances, behaviour or choices of the individuals concerned;
- have a prolonged or permanent impact on the data subject; or
- at its most extreme, lead to the exclusion or discrimination of individuals.

It is difficult to be precise about what would be considered sufficiently significant to meet the threshold, although the following decisions could fall into this category:

- decisions that affect someone's financial circumstances, such as their eligibility to credit;
- decisions that affect someone's access to health services;
- decisions that deny someone an employment opportunity or put them at a serious disadvantage;
- decisions that affect someone's access to education, for example university admissions.

Having laid out the scope of the prohibition, Article 22 continues with three exceptions:

2. Paragraph 1 shall not apply if the decision:
 - a) is necessary for entering into, or performance of, a contract between the data subject and a data controller;
 - b) is authorised by Union or Member State law to which the controller is subject and which also lays down suitable measures to safeguard the data subject's rights and freedoms and legitimate interests; or
 - c) is based on the data subject's explicit consent.

Here again, a number of questions can be raised. The reader is advised to carefully study the EDPB Guidelines on Automated Individual Decision Making and Profiling, to gain a proper understanding of how these exceptions must be interpreted.

⁴ Ibid. at 21.

⁵ Ibid. at 21–22.

1. In the cases referred to in points (a) and (c) of paragraph 2, the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.

So, in the case of a decision based on automated processing that is necessary for a contract or a decision based on consent, *access to human intervention is required*, both to express one's point of view and *to contest the decision*. This is related to recital (71), which adds another requirement:

In any case, such processing should be subject to suitable safeguards, which should include specific information to the data subject and the right to obtain human intervention, to express his or her point of view, to obtain an explanation of the decision reached after such assessment and to challenge the decision.

Here we find the right to obtain an *explanation of the decision*, which many authors interpret as being a precondition to be able to contest the decision (as required in Article 22.3). By now, a number of scientific papers have been published on 'the right to an explanation' and 'explainable AI', which are deemed highly relevant also due to potential unwarranted bias. This 'right to an explanation' can also be read into the transparency requirements in Articles 13.2(f), 14.2(g), and 15.1(h), which all require that the following information will be provided:

- the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases,
- meaningful information about the logic involved, as well as
- the significance and the envisaged consequences of such processing for the data subject.

Data controllers have a *legal obligation* to provide such information, both when the data has been provided by the data subject (Article 13), and when data has not been obtained from the data subject (Article 14), while data subjects have a *right* to obtain such information (Article 15). Note that the obligation to provide these three types of information does not depend on a request by the data subject but must be provided anyway. Just imagine what this could mean for an iot system that runs on real-time ML applications, or for online credit applications based on ML inferences of credit worthiness.

4. Decisions referred to in paragraph 2 shall not be based on special categories of personal data referred to in Article 9(1), unless point (a) or (g) of Article 9(2) applies and suitable measures to safeguard the data subject's rights and freedoms and legitimate interests are in place.

The exceptions generally do not apply to automated decisions that are based on Article 9 data. Now think of unintended machine bias based on proxies that result in indirect racial discrimination as described above in section 10.1. There is no case law yet on how this prohibition must be interpreted, but we can imagine that Article 22.4 may provide far-reaching protection if properly interpreted in a balanced way.

Article 22 repeatedly speaks of 'suitable measures to safeguard the data subject's rights and freedoms and legitimate interests.' The EDPB clarifies that this includes technical measures. They write:

Errors or bias in collected or shared data or an error or bias in the automated decision-making process can result in:

- incorrect classifications; and
- assessments based on imprecise projections; that
- impact negatively on individuals.

Controllers should carry out frequent assessments on the data sets they process to check for any bias, and develop ways to address any prejudicial elements, including any over-reliance on correlations.

Systems that audit algorithms and regular reviews of the accuracy and relevance of automated decision-making including profiling are other useful measures.

Controllers should introduce appropriate procedures and measures to prevent errors, inaccuracies or discrimination on the basis of special category data. These measures should be used on a cyclical basis; not only at the design stage, but also continuously, as the profiling is applied to individuals. The outcome of such testing should feed back into the system design.

These types of 'safeguards' exemplify how LPbD can be turned into an *operational requirement* that guides the design of personal data processing systems, ruling out unwarranted violations of data protection law, while providing practical and effective protection at the level of the technical and organizational infrastructure.

References

On machine learning

- Mitchell, Thomas. 1997. *Machine Learning*. 1st ed. New York: McGraw-Hill Education.
- Mitchell, Tom M. 2017. 'Key Ideas in Machine Learning'. In *Machine Learning*, draft for the 2nd ed., 1–11.

On p-hacking and other risks in ML

- Berman, Ron, Leonid Pekelis, Aisling Scott, and Christophe Van den Bulte. 2018. 'P-Hacking and False Discovery in A/B Testing'. SSRN Scholarly Paper ID 3204791. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=3204791>.
- Hildebrandt, Mireille. 2018. 'Preregistration of Machine Learning Research Design. Against P-Hacking'. In *Being Profiled: Cogitas Ergo Sum*. Amsterdam: Amsterdam University Press.
- Hofman, Jake M., Amit Sharma, and Duncan J. Watts. 2017. 'Prediction and Explanation in Social Systems'. *Science* 355 (6324): 486–88. <https://doi.org/10.1126/science.aal3856> (quotation at p. 487).

On bias in ML applications

- Angwin, Julia, Jeff Larson, Surya Mattu, and Kirchner. 2016. 'Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And It's Biased Against Blacks'. *ProPublica*. 23 May 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- Barocas, Solon, and Andrew D. Selbst. 2016. 'Big Data's Disparate Impact'. *California Law Review* 104: 671–732.
- Chouldechova, Alexandra. 2017. 'Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments'. *Big Data* 5 (2): 153–63. <https://doi.org/10.1089/big.2016.0047>.
- Yong, Ed. 2018. 'A Popular Algorithm is No Better at Predicting Crimes than Random People'. *The Atlantic*, January. <https://www.theatlantic.com/technology/archive/2018/01/equivant-compas-algorithm/550646/>.

On the potential and real effects of ML on public space, democracy, and freedom of expression

- Pariser, Eli. 2011. *The Filter Bubble: What the Internet is Hiding from You*. London: Penguin.

- Sunstein, Cass R. 2016. *The Ethics of Influence: Government in the Age of Behavioral Science*. Cambridge: Cambridge University Press.
- Tufekci, Zeynep. 2018. 'How Social Media Took Us from Tahrir Square to Donald Trump.' *MIT Technology Review*, September/October. <https://www.technologyreview.com/s/611806/how-social-media-took-us-from-tahrir-square-to-donald-trump/>.

Re the fundamentals of 'smart contracts'

- Buterin, Vitalik. 2014. 'A Next-Generation Smart Contract and Decentralized Application Platform. White Paper.' Ethereum Platform.
- Nakamoto, Satoshi. 2008. 'Bitcoin: A Peer-to-Peer Electronic Cash System.' <http://www.bitcoin.org/bitcoin.pdf>.
- Szabo, Nick. 1997. 'Formalizing and Securing Relationships on Public Networks.' *First Monday* 2 (9). <http://firstmonday.org/ojs/index.php/fm/article/view/548>.

Re the writing of 'smart contracts'

- Seijas, Pablo Lamela, Simon J. Thompson, and Darryl McAdams. 2016. 'Scripting Smart Contracts for Distributed Ledger Technology.' IACR Cryptology EPrint Archive 2016: 1156.

Re 'blockchain' and the GDPR

- Finck, Michèle. 2018. 'Blockchains and Data Protection in the European Union.' *European Data Protection Law Review* 4 (1): 17–35. <https://doi.org/10.21552/edpl/2018/1/6>.

Re compatibility of 'smart contracts' with Article 22 GDPR

- Art. 29 Working Party WP251rev.01, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679. https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053.
- CNIL, September 2018, 'Solutions for a responsible use of the blockchain in the context of personal data.' <https://www.cnil.fr/sites/default/files/atoms/files/blockchain.pdf>.

Re legal contracts and 'smart contracts'

- Allen, J.G. 2018. 'Wrapped and Stacked: "Smart Contracts" and the Interaction of Natural and Formal Language.' *European Review of Contract Law* 14 (4): 307–43. <https://doi.org/10.1515/ercl-2018-1023>.

- Cornell, N. and K. Werbach. 2017. 'Contracts Ex Machina'. *Duke Law Journal* 67 (2): 313–82.
- Raskin, M. 2017. 'The Law and Legality of Smart Contracts'. *Georgetown Law and Technology Review* 1 (2): 304–41.
- Verstraete, Mark. 2018. 'The Stakes of Smart Contracts'. SSRN Scholarly Paper ID 3178393. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=3178393>.

Re Legal by Design and Legal Protection by Design

- Filippi, Primavera De, and Samer Hassan. 2016. 'Blockchain Technology as a Regulatory Technology: From Code Is Law to Law Is Code'. *First Monday* 21 (12). <http://firstmonday.org/ojs/index.php/fm/article/view/7113>.
- Hildebrandt, Mireille. 2017. 'Saved by Design? The Case of Legal Protection by Design'. *NanoEthics*, August, 1–5. <https://doi.org/10.1007/s11569-017-0299-0>.
- Lippe, Paul, Daniel Martin Katz, and Dan Jackson. 2015. 'Legal by Design: A New Paradigm for Handling Complexity in Banking Regulation and Elsewhere in Law'. *Oregon Law Review* 93 (4). <http://papers.ssrn.com/abstract=2539315>.
- Van den Berg, Bibi, and Ronald E. Leenes. 2013. 'Abort, Retry, Fail: Scoping Techno-Regulation and Other Techno-Effects'. In *Human Law and Computer Law: Comparative Perspectives*, edited by Mireille Hildebrandt and Jeanne Gaakeer, 67–87. Ius Gentium: Comparative Perspectives on Law and Justice 25. Springer Netherlands. http://link.springer.com/chapter/10.1007/978-94-007-6314-2_4.

PART IV

FINALS

In the final part of this book law is confronted with two other types of normativity, those of ethics and those generated by computer code. Though this may seem an unnecessary extension of the work, it is pivotal to grasp where law ends and ethics starts, and where both law and ethics end and the force of computer code begins.

11

Closure: On Ethics, Code, and Law

This—final—chapter investigates the distinction between law, code, and ethics, their interrelationship and their interaction. It is a bonus chapter for those interested in the nexus of law and ethics, in the light of code- and data-driven information and communication infrastructures (ICIs).

In the introduction to Chapter 10 we have encountered MIT’s ‘moral machine’ thought experiment, which aimed to ‘mine’ opinions on the ethics of choices that self-driving cars may have to make.¹ I have qualified the experiment as befitting a ‘naive’ type of utilitarianism. In this chapter, I explain the assumptions that underlie the framing of the problem of ‘moral machines’ and discuss other traditional ways of framing ethical dilemmas. *This is necessary because they are part of our common sense and thus often serve as the hidden premises of ‘ethics in AI’ and similar attempts to ‘do good’ when developing code- or data-driven systems.* Such hidden assumptions play an important role even if one is not aware of them, and they must therefore be called out.

After providing the overview, I will clarify what differentiates law from ethics (section 11.1.6), as this is a book on law—not primarily on ethics. *Spoiler:* one of the main differences is that law provides *closure* whereas ethics remains in the realm of reflection as it does not have *force of law*. However, a second difference turns the previous statement inside out: whereas law and the Rule of Law introduce checks and balances and demand democratic participation (at least in constitutional democracies), ethics may be decided by tech developers or behind the closed doors of the board room of corporate business enterprise. It can thus obtain *the force of technology*.

This would imply that it is no longer law but also technology that provides closure, though not by way of democratically legitimated legislation. Instead, closure is provided by ethics, as embodied in the black box of R&D, the board rooms of Big Tech, and communities of developers that write and maintain

¹ <http://moralmachine.mit.edu>.

open source code or DLTs. Though the latter are not a black box for those knowledgeable on the technical side, they are black boxes for those who cannot read the code.

For a proper understanding of the role of ethics, code, and law in technology development we need to move beyond analytical distinctions. As demonstrated in Chapter 2, there is a special *relationship* between ethics and the Rule of Law, which implies that law and ethics *interact*. The example I will use throughout this chapter is not about the ethical dilemmas of driverless cars, but the question of algorithmic fairness (which obviously also regards decisions made by those who build the code for driverless cars). This will confront the force of law with the force of technology, requiring a new type of interaction between lawyers and computer scientists on how to ensure that ‘ethical design’ does not overrule the checks and balances of the rule of law. In that sense, some of the notions presented in Chapter 10 will resurface when discussing the relationship between code and law.

In the context of this chapter, I use the term ethics to refer both to *morality* (acting in a morally justified way) and to *moral philosophy* (inquiring into the types of moral justification one could develop). This also means that, for the purposes of this chapter, I use ‘ethical’ and ‘moral’ as synonymous.

11.1 Distinctions between Law, Code, and Ethics

Doing ethics can mean two different things:

1. being engaged in the philosophical subdiscipline of ethics; or
2. acting in a way that is ethical.

Though it may be tempting to invent an ethics for the onlife world, as if it does not matter what centuries of investigation into moral philosophy have brought us, this easily results in getting caught up in hidden assumptions. For instance, the MIT thought experiment is presented as if it has nothing to do with scholarly debates on the different schools of moral philosophy, but its framing of the problem rests on a specific variant of utilitarianism and incorporates a number of assumptions that are taken for granted without closer inspection.

To act ethically as an individual, one need not have studied ethics, but when reflecting on the ethical implications of, for example, bias in machine learning, it is crucial to take a step back before moving forward.

11.1.1 Utilitarianism and methodological individualism

Utilitarianism is focused on the consequences of our actions. For that reason, it is often equated with consequentialism. Utilitarianism is, however, a particular type of consequentialism, based on ‘methodological individualism’. This means that individual choices are assumed to be independent, such that collective choice is nothing other than the aggregate of individual choice. This is a highly contentious position, as individual choice is dependent on the anticipation of another’s choice and in part constituted by choice architectures that are in turn dependent on ICIs and informed by power relationships.

The interdependencies between individual and collective choice in complex systems such as human society are numerous and, in part, emergent. Simplifying them by assuming independent individual choice may be convenient from a computational point of view, but entirely inadequate as to real-world interaction. This is why rational choice theory may seem a nice tool to think about ethical choices, but it fails to register that as a tool it actually co-determines what it supposedly investigates. It creates a *framing problem*.

This relates to the second untenable assumption of ‘methodological individualism’, that is, that ‘means’ and ‘ends’ can not only be analytically distinguished (a very good idea) but ‘exist’ separately in our world (a highly problematic idea).

In the section on pragmatism (section 11.1.4), I will clarify the dependencies between means and ends as part of the framing problem that is inherent in any debate on ethics and AI. Though pragmatism also ‘thinks in terms of’ consequences, it does not assume the separation between means and ends that is assumed in utilitarianism.

For the sake of brevity, I discuss four intersecting types of utilitarianism, inevitably leaving many nuances aside: act- and -rule-utilitarianism, and maximum and average utilitarianism. All four emphasize that ethical choice must be made on the basis of the utility it generates. That is why utilitarianism feeds on cost-benefit assessments that in turn nourish a utilitarian calculus; it forms

the hidden assumption of risk assessment as a viable way to cope with the impact of new technologies. Because people may not agree on what constitutes utility, the consequences are usually discussed in terms of preferences or well-being rather than utility. That, however, raises the question of whether these preferences are given or framed, depending on the choice architecture presented by the environment. Well-being raises similar questions, because well-being is not necessarily an objective function of ethical choices (different individuals, groups, cultures, and societies may define well-being in contrasting and even incompatible ways). Therefore, I will stick to the concept of utility, taking note that it is the vanishing point of utilitarianism and in many ways a black box.

Act-utilitarianism says that the right act is the one that maximizes utility. The question obviously is ‘whose utility?’, because the maximization can be understood as an aggregate: the more peoples’ utility is satisfied, the better, or, as an average: the higher the average utility, the better. Many modulations are possible. Legal philosopher John Rawls might require that the outcome should at least optimize utility for the ‘least advantaged’, while still rewarding those whose actions increased the overall scope for utility. This is coined the *maximin principle* and will be explained below under deontological reasoning (section 11.1.2), and under justice, legal certainty, and instrumentality (section 11.2.1), since Rawls is not a hard-core utilitarian. The most important point here is that in act-utilitarianism each act is isolated as if it were a stand-alone ‘device’. The moral machine experiment required visitors to provide a moral preference based on limited access to context, background, and circumstances—as if the situations occur in a vacuum.

Rule-utilitarianism was meant to resolve the problems generated by act-utilitarianism. It basically proclaims that the right act is the one that aligns with a rule that would—if everybody were to follow it—achieve maximum utility. As with act-utilitarianism, some may prefer average to maximum utility, or follow Rawls’ maximin principle. Rule-utilitarianism shares with act-utilitarianism the assumptions of ‘methodological individualism’ and the separation of means and ends. This results in a propensity to quantify the problem by way of game theory (assuming rational agents) or behavioural economics (assuming that though human agents may be irrational, they are nevertheless predictable as to their irrational behaviours).

I can now explain why I believe that MIT’s ‘moral machine’ experiment rests on a ‘naive’ type of utilitarianism. Either it aims to unearth the moral preferences of website visitors as to the desirable consequences of a series of particular acts, in which case all the problematic assumptions

of act-utilitarianism apply. Or it aims to uncover the *moral preferences* of website visitors as to the type of rules that should inform the behaviour of autonomous vehicles, with regard to a specified act. In that case act-utilitarianism is conflated with rule-utilitarianism, because the whole idea of rule-utilitarianism is to achieve guidance at a higher level of abstraction (not case-based but rule-based).

The researchers could object that their study is just an objective data-driven investigation into the moral preferences of 40 million webvisitors, and should not be confused with an ethical inquiry. They might assert that the study does not endorse any theory of ethics and does not contain any bias towards utilitarianism. Philosopher of science Karl Popper would respond that cognition and even perception is not possible without an underlying theory that frames the issues under investigation. In this case, the methodological individualism that underpins utilitarianism clearly frames the experiment and configures the kind of choices webvisitors are presented with. These choices are then qualified as their *given preferences*, and treated as *independent variables* that can be correlated with, for example, 'cultural traits', 'economic predictors', and 'geographical proximity'. As Michel Callon and John Law wrote, quantification (numerical data) is necessarily preceded by qualification (grouping specific instances under the same heading of a specific variable or feature). Though there is nothing wrong with such qualification, we need to become aware of the definitional choices they imply, and the framing issues they generate. Below I will give an example of assessing algorithmic fairness in a way that calls out these choices and shows some of their implications (section 11.1.5).

Here, it suffices to highlight that both types of utilitarianism would ultimately require a way to measure and maybe even weigh preferences (would, e.g. a preference to save white folk over coloured folk 'count' at all?). Usually, these kinds of preferences are *agent-dependent*, because my choice for a behavioural rule or action may depend on whether I am in the car or outside. It is entirely unclear how webvisitors developed their preferences, which makes the whole experiment a rather hazardous attempt to contribute to an informed debate on the ethics of self-driving cars. To seriously understand ethically relevant preferences, we should impose a *veil of ignorance*, requiring us to decide without knowing whether we will be the victim or not. However, this may bring rule-utilitarianism rather close to deontological imperatives, since the reasons that inform my agent-independent choice may differ from those that inform my agent-dependent choice, which introduces a moral criterion that is not part of the notions of either utility, act, or rule.

Let us now turn to algorithmic fairness, inquiring how it would fare under various types of utilitarianism. The problem is that neither maximum nor average utility would solve the problem of the disparate impact of various types of bias in machine learning. In the aggregate, unfair bias may increase utility (whether maximized or on average), but some categories of individual persons may find that their preferences are ignored or diminished. Clearly, fairness is a moral criterion that cannot easily be fitted into the logic of either act- or rule-utilitarianism.

11.1.2 Deontological reasoning: respect for human autonomy

Deontological reasoning is about people doing the right thing for the right reason, without taking into account the effects. Deontological reasoning is about duties, not about consequences, and can be traced back to Kant's categorical imperative. Kant distinguished between a *hypothetical imperative*, which makes a decision depend on the consequences it is expected to generate (often assessed from the perspective of one's *personal interest*), and the *categorical imperative*, which makes a decision depend on the moral justification it involves (notably respecting the *autonomy of others*).

Kant formulated different versions of the categorical imperative. I am quoting them here from the renowned *Stanford Encyclopedia of Philosophy*, to give the reader a taste of the complexities that deontological reasoning may involve, making it seemingly less amenable to computational translation than a utilitarian calculus (though the problem of defining utility creates the same kinds of problems). Note that the emphasis on individual moral autonomy does not depend on the methodological individualism of utilitarianism, as the maxims to be discussed do not depend on an aggregate utility, but on the extent to which a maxim implies that everyone's autonomy is respected.

1. act only in accordance with that maxim through which you can at the same time will that it become a universal law.

According to the *Stanford Encyclopedia of Philosophy* this implies:

First, formulate a maxim that enshrines your reason for acting as you propose.

Second, recast that maxim as a universal law of nature governing all rational agents, and so as holding that all must, by natural law, act as you yourself propose to act in these circumstances.

Third, consider whether your maxim is even conceivable in a world governed by this law of nature. If it is, then,

fourth, ask yourself whether you would, or could, rationally will to act on your maxim in such a world.

If you could, then your action is morally permissible.

2. we should never act in such a way that we treat humanity, whether in ourselves or in others, as a means only but always as an end in itself.

According to the *Stanford Encyclopedia of Philosophy* this implies:

First, the Humanity Formula does not rule out using people as means to our ends.

Second, it is not human beings per se but the 'humanity' in human beings that we must treat as an end in itself.

Third, the idea of an end has three senses for Kant, two positive senses and a negative sense.

Finally, Kant's Humanity Formula requires 'respect' for the humanity in persons.

3. the Idea of the will of every rational being as a will that legislates universal law.

According to the *Stanford Encyclopedia of Philosophy* this implies:

in this case we focus on our status as universal law givers rather than universal law followers.

This is of course the source of the very dignity of humanity Kant speaks of in the second formulation.

A rational will that is merely bound by universal laws could act accordingly from natural and non-moral motives, such as self-interest.

But in order to be a legislator of universal laws, such contingent motives, motives that rational agents such as ourselves may or may not have, must be set aside.

4. act in accordance with the maxims of a member giving universal laws for a merely possible kingdom of ends.

According to the *Stanford Encyclopedia of Philosophy* this implies:

it requires that we conform our actions to the laws of an ideal moral legislature,

that this legislature lays down universal laws, binding all rational wills including our own, and

that those laws are of ‘a merely possible kingdom’ each of whose members equally possesses this status as legislator of universal laws, and hence must be treated always as an end in itself. The intuitive idea behind this formulation is that our fundamental moral obligation is to act only on principles which could earn acceptance by a community of fully rational agents each of whom have an equal share in legislating these principles for their community.

To sum up, this type of deontological reasoning is grounded in a fundamental respect for the autonomy of each person, requiring us to act according to rules that any person could accept as the right rule.

Notice that ‘could’ is not equivalent with ‘would’, because ‘would’ may depend on self-interest, whereas ‘could’ depends on valid moral reasons to agree on the rule, taking into account other persons’ autonomy. This abstracts from personal preferences and from mere *acceptance* of rules, demanding that rules are instead *acceptable* from the perspective of a *rational universal consensus* on how each person’s autonomy is best respected. This entails a reconstructive morality in the sense that one’s actions should be justifiable as fitting a general rule that anybody would agree to *behind a veil of ignorance* (not knowing what would be in one’s personal interest, thus turning the above ‘could’ into a ‘would’).

Clearly, the assumption of a rational universal consensus is problematic, not because people have different interests (the veil of ignorance solves that problem), but because people have different ideas about the value of such interests and about their ranking (e.g. preferring community over liberty, or equality over community). We shall return to this when discussing pragmatism.

How would algorithmic fairness fit with the framework of deontological reasoning? One way to approach this would be to ask whether bias in algorithmic decision-making systems violates the autonomy of some human agents, while respecting the autonomy of others. The inequality goes to the heart of the matter, since the categorical imperative does not allow more or less respect for a person’s autonomy; either it is respected, or it is not respected. From the perspective of Kant, autonomy is not respected if there is no universal rule that justifies disparate treatment. To assess whether this is the case we need to ask whether different treatment *would be consented to if one had no idea whether one would benefit or lose out due the algorithm.*

This thought experiment was proposed by John Rawls under the heading of ‘the veil of ignorance’.

This veil of ignorance inspired Rawls' ethical *maximin principle* that explains under what conditions inequality is not unfair. Imagine there is one cake, to be shared by a group of people. Some of them may figure out ways to enlarge the cake. Since one is behind the veil of ignorance, there is no way of knowing whether one is amongst those who could 'grow' a bigger cake or not. The maximin principle says that by default everyone should obtain an equal share of the cake. However, it would be fair that those who manage to enlarge the cake, should be given a larger share than the others. This should, nevertheless, not result in those with the smallest parts to end up with even less than before. In fact, they should benefit from the enlargement of the cake, though not to the same extent as those who made it grow. This way, those who increased the shared cake are rewarded for their contribution (just desert), while taking care that the least advantaged share in the increase (fair distribution).

Rawls basically combines two types of justice as fairness in his maximin principle: *distributive and corrective justice*. We will return to this when discussing justice (section 11.2.1).

There may be a preliminary matter that is even more to the point here: can the automated application of an algorithm ever be respectful of the autonomy of those subject to its decisions? Could it be that algorithms necessarily use people only as a means and cannot ever respect their autonomy, due to the nature of machinic decision-making? This is a pivotal question and I believe that the answer depends on a number of factors that relate to the extent to which human oversight and human intervention are ruled out. I would not categorically reject algorithmic decision-making, because one can argue that abstaining from its usage could result in invisible unfair treatment by human beings (whether deliberate or unintended). One could argue—in that case—that abstaining from algorithmic decision-making shows disrespect for the autonomy of those subject to the decision.

11.1.3 Virtue ethics: perceiving the good and doing what is right

Rule-utilitarianism and deontological reasoning based on the categorial imperative seek ethical guidance in abstract rules that should be applicable independent of the personal characteristics or inclinations of the acting agent. Virtue ethics is less impressed with abstract justification, as it is focused on the moral character developed by the actor. This is not a matter of agent-dependent reasoning based on the self-interest of the agent, but a matter of

highlighting the need for individual agents to practice and develop their moral compass. The idea is that human agents are not born with such a compass, but need to gain experience in real-life situations, building what Aristotle called *phronesis* or practical wisdom. In the context of virtue ethics, the point is not to submit oneself to abstract rules but to elicit the right rule for the situation at hand. This is a matter of acuity and judgment rather than the application of existing rules or a calculation of utility.

Where utilitarianism and deontological ethics are focused on reasoning about the right decision when facing contradictory duties or conflicts of interest, virtue ethics is about the *perception of what is good and acting on it*.

As Varela wrote in his work on *Ethical Wisdom*:

As a first approximation, let me say that a wise (or virtuous) person is one who knows what is good and spontaneously does it. It is this immediacy of perception and action which we want to examine critically. This approach stands in stark contrast to the usual way of investigating ethical behavior, which begins by analyzing the intentional content of an act and ends by evaluating the rationality of particular moral judgments.

Aristotle distinguished between two types of knowledge: *theoretical knowledge* or *episteme*, and *practical wisdom* or *phronesis*.

Whereas *episteme*, according to Aristotle, is a matter of reasoning and theoretical insight, *phronesis* is a matter of experience, action, and perception. Young men (Aristotle was not interested in women) are great in achieving *epistemic* knowledge, whereas *phronesis* can only be achieved in the course of a lifetime. Perhaps virtue ethics is the most interesting type of ethics in an onlife world, where non-human agents challenge our understanding of moral agency. It seems clear that machines may develop something akin to epistemic knowledge. They will, however, by definition be excluded from developing virtues or practical wisdom. This is related to the difference between knowledge and wisdom, and between rationality and moral character. Wisdom and moral character require a type of acuity that implies both ambiguity and good intentions, together with skilled intuition, a kind of tacit knowledge that incorporates virtues such as prudence, temperance, courage, and justice. It is hard to imagine that a deep learning algorithm develops any of these characteristics in its relationship with other agents, even if it beats grand masters in chess, Go, and whichever other closed game with well-defined rules.

How would algorithmic fairness fare with virtue ethics? Could one define the virtue of justice such that it can be formalized and computed? Might Aristoteles' distinction between distributive and corrective justice (section 2.2.2) lend itself to research designs that detect unfair bias, while repairing whatever bug led to the violation of justice?

It seems that virtue ethics is based on a specific type of *incomputability*, notably regarding the relational nature of human agency and human intercourse, thus confirming that fairness cannot be calculated (though it can—paradoxically—be framed and calculated in many ways none of which can claim to have the one right answer). This may indicate that the concept of an ethical algorithm is an oxymoron that ignores the *undecidability* of virtuous action and fair decision-making. Not because humans are more often right than machines, but because the relational nature of virtuous action has no place in a system that can only ever execute code (whether in the form of deterministic self-executing code or in the form of inductive inference engines).

11.1.4 Pragmatist ethics: taking into account

The founding father of pragmatism, Charles Saunders Peirce, developed the so-called 'pragmatist maxim':

Consider what effects, which might conceivably have practical bearings, we conceive the object of our conception to have. Then, our conception of those effects is the whole of our conception of the object.

It should be clear that pragmatism is deeply *consequentialist*, to the extent that the meaning of the words we use is defined in terms of the anticipated effects of their usage. This leads pragmatism, in the end, to the acknowledgement that *means co-determine or reconfigure ends* in a way that makes their separation a naive though sometimes productive thought experiment (in philosophical terms this implies that means and ends can be analytically distinguished but not ontologically separated).

This clearly has implications for ethics, as it highlights that the way we try to achieve our objectives shapes them, also in the realm of ethics. In the context of utilitarianism, technologies are often seen as neutral tools, ignoring the way they enable and constrain both intended and unintended effects. In the context of deontological ethics all that seems to matter is one's moral duties to other agents, based on an abstract rational consensus that fails to take into

account the situatedness of human agency. This results in moral duties that abstract from the mundane means of executing them, thus missing out on their impact on human autonomy. Other than Kant, an ethical pragmatist would not assume or postulate an autonomous human subject, but seek to uncover the real-life conditions for autonomous agency.

Virtue ethics seems highly relevant in the realm of value-sensitive design, as the success of ‘ethical design’ will depend on the skills needed to make value-sensitive design work. But it is pragmatism that has the clearest understanding of the normative implications of designing a technology one way or another, precisely because it is already aware of how the means shape the goals. A pragmatist ethics shares awareness of the situatedness of the human agent with virtue ethics, and a sensitivity to the importance of experience, since pragmatism highlights the need to anticipate consequences (albeit not in the utilitarian sense). As with virtue ethics and utilitarianism, a pragmatist ethics is less impressed with the universal moral duties of deontological reasoning, and it endorses a more situated understanding of human autonomy.

From a pragmatist perspective, algorithmic fairness is clearly an ethical concern, since pragmatism acknowledges that any technology that is used as a tool to achieve some specific goal will:

1. result in what is usually called side-effects,
2. redefine the goal in terms of the means to achieve it,
3. thus reconfiguring the affordances of the environment of the human agent(s),
4. which will probably have *normative effects* that may require a *moral assessment*.

We can point to the work of Helen Nissenbaum, notably to her ‘contextual integrity’ (CI) heuristic, that traces the implications of novel types of technologies, providing a step-by-step assessment of how the environment is changed and how this may affect the legitimate, context-based expectations of human agents. One of the consequences of introducing novel technologies may be a redistribution of risks and benefits within and across contexts, which may reinforce existing inequalities or even create new types of inequality. Her analysis fits with the core assumptions of a pragmatist ethics, it moves beyond privacy and provides a coherent framework to assess fairness as an ethical value that may be disrupted.

Note that contextual integrity does not equate fairness with equality. As we have seen above, when discussing Rawls’ maximin principle, treating different people equally may actually be unfair. Think of Anatole France’s famous finding that: ‘In its majestic equality, the law forbids rich and poor alike to sleep under bridges, beg in the streets and steal loaves of bread.’ The balance

that must be struck between corrective and distributive equality requires choices that assume a moral and a political evaluation of what counts as fair under what conditions. There may be clear indications of unfair treatment, but it is not easy to come to an agreement on what constitutes fair treatment.

Ultimately, this is a *moral choice* that individuals and societies will have to decide on, and a *political choice*, for instance, to enact legal norms that prohibit certain actions as unfair and therefore unlawful.

11.1.5 The difference that makes a difference: closure

Before drawing conclusions regarding the major differences between law, code, and ethics, I will present the reader with an excerpt of a blogpost on *Medium* by the Berkman Klein Centre at Harvard University, on the so-called ‘Detain/Release’ teaching module. This module simulates pre-trial court decisions on whether to detain or release a defendant based on the available assessment of recidivism risk:

We wanted students to put themselves in the role of a judge, and think about how they would make pretrial detention decisions. We began with a tutorial run that students completed on their own: ten defendants, no risk assessments.

After that, we divided students into groups and had them do three full runs of the simulation. We wanted students to talk about how they made their decisions, during and after the simulation runs. By the third run, we found that students are invested in the simulation and in the detention and release decisions they’ve made.

Throughout, we were deliberately opaque about how the simulation worked—about how accurate the risk assessments actually were, and about what probabilities ‘low’, ‘medium’, and ‘high’ corresponded to. For the most part, no one asked, either in our classroom or during our tests of the simulation.

Despite that, as they progressed through the lesson, students began to feel more confident and assured in their detention and release decisions. They built interpretive systems to quickly make decisions from the information they had been given. Some of their rules and systems were expected: high violence usually meant detention. Others, less so: after seeing two female defendants fail to appear, one team began detaining women by default.

After the third and final run, we showed students the consequences of their decisions, with one last dashboard view: How did pretrial detention decisions affect defendant outcomes?

You detained 159 people.

68 of those people will plead guilty before trial. Of these, 37 will have done so as a result of being detained.

Of the people who go trial, 68 will be found guilty. 20 of them will have been convicted because their detention inhibited their ability to mount a defense.

Upon their release, 37 of the people you detained will face eviction. They will struggle to find housing. 99 will continue to be out of work three years after they are released.

In total, the people you detain will spend an average of 200 days waiting for their cases to be disposed.

The final dashboard view: consequences.

This reveal takes the air out of the room. It drives home the framing power of the risk assessment tool we had presented them: students relied on it, deeply, despite receiving no promises about its accuracy, and ‘corrected’ for it in random ways. This had consequences.

The aim here is not to take sides on who are right or wrong with regard to the use of pretrial software to conduct a risk assessment, or on whether human judges do better than the software. The point here is to demonstrate that MIT’s thought experiment will only contribute to a sustained reflection on e.g. algorithmic fairness if the framing problem is faced and addressed. The Berkman Klein module on the ‘Detain/Release’ simulation nicely shows how software systems can lure decision-makers into accepting assumptions and implications that should be called out before being put into action.

What can we learn from the above on the difference between law, code, and ethics?

1. The study of ethics concerns a reflection on the justification (whether utilitarian or deontological) of decision-making that affects human agents and human societies, and/or the development of practical wisdom (virtue ethics), and/or the study of how the means to reach desirable goals reconfigure those goals as well as the values they incorporate (pragmatist ethics). The study of ethics and the development of practical wisdom do not have the force of law; they do not (and should not) provide closure on how to act or how to design our ICIs.
2. Positive modern law provides closure in a way that ethics cannot and should not do, since a constitutional democracy rules out the imposition of a specific ethical stance. Precisely because ‘we’ do not agree on ethics, we need

law to coordinate our behaviour in a way that provides for legal certainty and justice—in a way that sustains the instrumental role of the law (section 2.2.2). The closure of modern law is directly related to its ‘positiveness’ (it is enacted—posited—by the legislature, its interpretation is decided by independent courts, whose verdicts are enforceable due to the monopoly of violence). The fact that law provides closure does not, however, imply that there is no relationship whatsoever between law and ethics. The fundamental requirement of justice forms the interface with ethics and determines the inner morality of the rule of law, which is a specific type of meta-ethics. We shall return to this in the next section (section 11.2.1).

3. Acting ethically concerns making the right decision, both at the level of individual choice and at the level of designing the legal, political, and technical *choice architectures that frame such choice*. Both types of decisions interact, and they achieve closure to the extent that they foreclose the effects that another decision might have generated. In the case of design choices, the impact may be substantial.
4. The development and implementation of computer code in a variety of algorithmic decision-making systems may achieve closure, due to the choice architectures they present. At this moment, such closure is not part of democratic participation and there is no way to guarantee that the checks and balances of the rule of law are integrated.

One could conclude that, whereas ethics is not a competitor of law, algorithmic decision-making systems are just that.

11.2 The Conceptual Relationship between Law, Code, and Ethics

Ethics is both more and less than law: it is more because many ethical concerns are not addressed by the law and less because the outcome of ethical considerations are not necessarily transformed into legal norms and thus not enforceable by way of law. As indicated above, since we often do not agree on ethical rules, values, or choices, the law mainly integrates ethical principles and considerations at a meta-level—for example, to make sure that ethical choice is not systematically overruled by economic interest. The idea is that law and especially the rule of law creates space to develop one’s practical wisdom and to act in accordance with the kind of rules one believes everyone should follow (seen from behind a veil of ignorance).

I will now first return to section 2.2.2 to clarify once again the relationship between law and ethics at the level of law's foundational architecture. After that I will flesh out how this foundational architecture relates to the employment of computer code when making legally relevant decisions.

11.2.1 Justice, legal certainty, and instrumentality

The goals of ethics can be summed up as 'acting in the right way', which assumes having taken the right decisions, taking note that these decisions may be implicit in our actions since much of our ethical knowledge is tacit and hard to spell out. The study of ethics hopes to explain how our actions can be justified, by, for example, referring to values such as liberty, equality, and autonomy. Though part of moral philosophy assumes that a universal rational consensus about what constitutes a right action is possible, the problem with ethics is precisely that there is no such consensus (neither is there a consensus that we should try to reason towards such a universal rational consensus). In point of fact, constitutional democracies take the position that *it would be unethical to impose the ethics of a majority on minorities*, let alone that the ethics of a minority should reign over a majority. But, as some would remind us, this position itself is precisely the kind of universal rule we need in a meta-ethical framework.

Law cannot disentangle itself completely from ethics. On the contrary, law and the rule of law embrace a pragmatic *meta-ethics* that integrates a system of institutional checks and balances that safeguard the freedom to live according to one's own ethics—though within the limits needed to guarantee equivalent safeguards for others. This means that law is concerned with a specific type of justice, closely aligned but not equivalent with legal certainty. As discussed in section 2.2.2, law has to serve three different, partly overlapping and often incompatible goals: those of justice, legal certainty, and instrumentality.

Justice concerns the combination of distributive and corrective justice that ensures that the law:

1. treats similar cases equally to the extent of their similarity; and
2. provides for just desert in proportion to whatever elicits the desert (e.g. committing a tort or a criminal offence or creating added value for society).

Though we can agree that people should be treated equally, we may not always agree on *what counts as equal* and we must also admit that treating everyone equally badly does not agree with our sense of justice, because it cannot be that this is deserved.

Above, in section 11.1.2, we discussed Rawls' maximin principle as a way to combine both types of justice, under the heading of 'justice as fairness'. Even in that case, we need to take a series of decisions about how this balance can or should be struck, leaving room for choice, interpretation, and contestation.

In the end, political decisions must be made, for example, about what constitutes a fair market, enacting the relevant legislation, followed by legal decisions that apply what the legislature enacted. From that moment onwards, the law will take over and make sure that law's *instrumentality* in terms of policy goals set by the legislature is achieved in alignment with *legal certainty* (foreseeability) and *justice* (distributive and proportional equality). Here again, courts will have to take decisions on what counts as equal and what is deserved. Sometimes, a decision may be fair but unforeseeable, foreseeable but unfair, or it may resist instrumentality to safeguard foreseeability or violate fairness to assure instrumentality.

There is no way to resolve—at an abstract level—the tension between the three goals of the law: justice, legal certainty, and instrumentality. What matters is that any and all legal decision(s) must be justifiable as striving to serve all three goals, thus *sustaining rather than resolving* the tension between them. This 'demand' can be termed a *meta-ethics* that basically enables people to develop their own moral competences. For instance, if ethical values such as privacy and fairness are left to 'the market', companies that build their systems in accordance with these values may be pushed out of the market (because they have to make costs that other companies externalize). If, however, the law puts a threshold in the market by requiring and enforcing companies to integrate these values into their systems, companies can 'afford' to act ethically.

11.2.2 Law, code, and the rule of law

In the previous subsection, we have seen that the relationship between law and ethics can be traced to the fact that ethics informs a rule of law that:

1. requires that the *instrumentality* of law as a means to achieve goals set by the legislature,

2. is constraint by both the foreseeability and stability of the law and its equal application (*legal certainty*),
3. based on the idea that governments must demonstrate equal respect and concern for all citizens (*justice*).

Though justice is an ethical value, its role in law is limited by the instrumentality of the law (an orientation towards goals defined by the legislature, or, in the case of contract, by contracting parties) and by the demands of legal certainty (the ‘positivity’ of the law, meant to ensure both the enforceability of the law and the integrity of the law as a whole). This confirms that law is both more and less than ethics.

This raises the question of how law and the rule of law relate to code, an issue already addressed in Chapter 10, notably section 10.3 where we distinguished ‘legal by design’ from ‘legal protection by design’. Here we look more broadly at algorithmic decision-making systems, whether in the private or the public sector, without focusing on systems that supposedly execute legal norms.

What if computer code is employed to decide individual cases for reasons of effectiveness, expediency, and scale? How does this relate to law and the rule of law and to ethics?

1. First, as discussed above, algorithmic decision-making changes the relationship between law and ethics to the extent that ethical choices may gain the force of technology, thus *becoming a competitor of law* in terms of enforceability.
2. Second, though both types of enforceability have a fundamentally different nature, they both affect those subject to their decisions, potentially *reducing the space for ethical choice*.

Technological enforcement reduces the space for ethical choice, because ethical choice assumes the freedom to act otherwise and room to develop alternative ethical positions. The space for ethical choice can be occupied either by legal obligations or by computer code. Insofar as legal norms impose particular ethical choices, the relevant conduct is turned into legal compliance. The same can be said about computer code that forces ethical choices upon people or companies, since—in that case—the choices are no longer made by those people or those companies.

The difference between law and computer code, however, is that a legal norm can in principle be *disobeyed*, whereas code that manages to constrain the behavioural options of people or companies may not leave any room for disobedience. This is a significant difference between law and technology, meaning that law leaves room for ethical choices even where it imposes its norms (think of civil disobedience), whereas computer code may leave no such room. Think of an algorithm that automatically allows advertisers to target white men for higher paid jobs, thus excluding women and people of colour from being informed about these jobs. The ethical choice that is at stake here is the choice of, for example, a website owner to disallow this type of unfair targeting. Since the algorithm is trained to increase ad revenue it may be difficult if not impossible to root out this type of algorithmic output, to the extent that the algorithm ‘finds’ that such exclusionary targeting increases ad revenue.

But we can go a step further: what if we could develop a *meta-algorithm* that puts constraints on this type of algorithms, ensuring they will necessarily be fair. What if we can develop an ‘ethical algorithm’, based on the formalization of a specified concept of fairness? Though this may be a wonderful way to achieve a specified type of fairness, it will reduce or transform the space for ethical action. Perhaps, in this case, the space for ethical action is restricted to those who understand the code and/or to those who can decide on the employment and the development of the code.

The reduction of the space for ethical choice will necessarily result in a loss of space to practice one’s moral compass. As Roger Brownsword has argued, this also goes for the law. If we develop algorithms that are ‘legal by design’ or ‘ethical by design’, we diminish the space of law or ethics in favour of ‘technological management’. This may ultimately impact our understanding of ethics and law, notably where some may argue that the technological management of our choice architectures is a better way to achieve a ‘good’ society than either law or ethics.

11.3 The Interaction between Law, Code, and Ethics

By exploring the distinctions between law, code, and ethics, and their relationship, we have prepared the ground for a study of their interaction. At a

conceptual level, I will do this by discussing ‘by design’ approaches to law and ethics, and, at a more concrete level, I will do this by determining how law and ethics interact with code in the context of algorithmic fairness.

11.3.1 ‘By design’ approaches in law and ethics

In section 2.1 I wrote that ‘[l]egal certainty, one of the core values of the law, is not about fixating the meaning of legal norms once and for all. Instead, legal certainty targets the delicate balance between stable expectations and the ability to reconfigure or contest them.’

This implies that legal certainty *resists formalization*, since this would freeze the meaning of legal norms, reduce their adaptive nature, and diminish their contestability (only those who understand the code can contest it).

This similarly goes for ethics, which may be even more adaptive, as it is not constrained by the requirement of legal certainty and closure.

Code, however, *implies formalization*, it cannot exist without an act of translation that removes ambiguity and defines in precise and increasingly machinic terms what problem is being solved (from source code through the compiler to programming language or object code). Formalization removes the elasticity and adaptiveness that is inherent in human language.

Recall the pragmatist definition of meaning (section 11.1.4):

Consider what effects, which might conceivably have practical bearings, we conceive the object of our conception to have. Then, our conception of those effects is the whole of our conception of the object.

This definition is particularly apt for understanding what language ‘does’, because it highlights *the anticipatory nature of language usage* and the meaning it generates. In section 2.1.2 I briefly discussed *speech act theory* when explaining the *performative character of the law*; if specific legal conditions are fulfilled, law attributes specified legal effects. For instance, the meaning of ‘murder’ is defined by a combination of legal conditions that generate the legal effect of some action ‘counting’ as murder. This means that whoever performed this action becomes punishable.

Computer code is capable of similar operations, though here we are not discussing ‘effects, which might conceivably have practical bearings’ but a preconceived and determined set of effects (even if the complexity is such that we—due to our bounded rationality—cannot foresee them all).

Code does not produce meaning but ‘mere’ effects, at the level of its integrated circuits, its logical operations, and decisional throughput and output (including effects in the real world as, e.g. in an internet of things (iot), or when using fintech, search engines, or social networks). Many of these effects may not only be *unforeseen* but also *unintended*, especially where the output pours out into the real world. This is where ‘by design’ approaches in law and ethics become interesting, in part because these limitations may also apply to ‘by design’ approaches that rely on adapting code as a solution.

Privacy by design has long been an example of a ‘by design’ approach in ethics, because there was no legal obligation to integrate privacy at the level of design. *Data protection by design (DPbD)* is an example of a ‘by design’ approach in law, at least within the jurisdiction of the General Data Protection Regulation (GDPR), because since 2018 this is a legal obligation (section 5.5.2.9).

This has implications for both privacy and other fundamental rights, for example, the right to non-discrimination:

- First, one may want to counter existing *privacy problems* by defining them in a way that lends itself to formalization and then figuring out a way to resolve the problems as defined. For instance, k-anonymity and differential privacy define privacy in terms of the hiding of data and/or the non-identifiability of data in aggregate data or in the patterns inferred from it. Based on that definition, one can develop metrics that enable one to prove mathematically to what extent privacy is protected. One could, for example, claim that differential privacy better protects privacy than k-anonymity, while still retaining aggregate data and inferred information that serves its purpose.
- Similar attempts to counter the undesirable implications of algorithmic decision-making systems are being made with regard to *fairness*. The problem is defined in a way that allows for formalization and is subsequently resolved—at that level—with regard to that specified definition of (un)fairness. To the extent that unfair treatment is unlawful, the legal requirement of DPbD may require that algorithmic decision-systems are designed in ways

that mitigate the unfairness, because DPbD is not limited to privacy. As discussed in section 5.5.2.9, Article 25 GDPR defines DPbD with regard to ‘risks of varying likelihood and severity for rights and freedoms of natural persons’. The fundamental right to non-discrimination (e.g. Article 21 of the Charter of Fundamental Rights of the European Union (CFREU)) thus requires a ‘by design’ approach *in law* regarding a lack of fairness that violates the right to non-discrimination.

However, this right is limited to discrimination based on a specific type of grounds (Article 21 CFREU speaks of any ground such as ‘sex, race, colour, ethnic or social origin, genetic features, language, religion or belief, political or any other opinion, membership of a national minority, property, birth, disability, age or sexual orientation’), and may be justifiable if specific conditions apply (e.g. reserving the payment of a pension to people beyond a certain age, reserving pregnancy leave to women, and reserving positive discrimination to a disadvantaged minority).

To the extent that algorithmic decision-making systems result in violations of fairness that is not unlawful in terms of DPbD, the obligation does not apply. In that case, a design approach could be based on ethical considerations. In the next subsection I will discuss fair computing as an example of ‘fairness by design’ that may in part be warranted under the legal obligation of DPbD and in part be based on a ‘by design’ approach to ethical issues around fairness in computing.

11.3.2 Fairness by design and ‘fair computing’ paradigms

Before heading into ‘fairness by design’ I need to address two preliminary issues.

- In the first place it is crucial to acknowledge that the formalization of a problem may—unintentionally—result in misrepresenting the problem. It may be that some forms of unfairness can be detected, whereas others remain elusive. The temptation may be to address what can be defined and resolved, whereas the problem that really bugs people resists the kind of *generalization* that is implied in *formalization*. This is an issue that must be squarely faced on pain of wasting time, money, and effort on a kind of technological solutionism that is not informed by real world problems.

Our tacit knowledge of what is unfair may be difficult to retrieve in more explicit expressions that may be both over- and underinclusive; tacit knowledge may be too complex to render in propositional terms without losing several dimensions that make a difference. This may even be due to the fact that we may have no words to describe our perception of injustice, resulting in what Miranda Fricker has coined ‘hermeneutical injustice’. This, in turn, is related to the fact that problems of fairness require framing, and different ethical positions will result in different framings. So, whereas some may find price-discrimination unfair for those who pay a higher price, others will argue that this is actually beneficial for those who have less to spend as it can lower their price. In reality, the higher price may, however, be paid by those with lesser means. For instance, health insurance may be more expensive in neighbourhoods with more low-income residents as statistically they have more health problems. Some will find this justified, from the perspective of the insurance company, others will find this unjustifiable, based on a Rawlsian veil of ignorance.

- The second issue that must be faced is that technical solutions may be used to legitimize algorithmic decision-systems that are fair in one particular way, but otherwise massively invasive and perhaps unfair in many other ways. As Powles and Nissenbaum have argued, providing this type of solution may distract attention from the preliminary question whether we want to actually replace human judgment with computational decision-making, in domains such as medicine, accounting, law, or education. These questions should not be asked at a high level of abstraction, but addressed in concrete situations, taking into account *how* the introduction of algorithmic decision-making may impact our information eco-system, the distribution of risk, and the capabilities of the human beings that will suffer or enjoy the consequences.

Having drawn attention to these preliminary issues, I believe that it is nevertheless pivotal to invest in researching and exploring ‘fairness by design’. Section 10.1.2 has provided an analysis of discrimination in parole decisions that are based on proprietary software, demonstrating that different people and organizations frame the issue of fairness differently, ending up in a dead-lock between those who claim statistical objectivity and those who argue that individual persons are in point of fact wrongly discriminated against, due to aggregate profiles that do not apply to them (the fact that 87 per cent of black people recidivize does not mean that every black person has a chance of 87 per cent to recidivize). Here we see the crucial difference between (1) ethical notions of unfairness that are by definition contestable; (2) legal notions of unfairness that are reasonably circumscribed but remain contestable on legal

grounds; and (3) computational notions of unfairness that are necessarily disambiguated to cater to the need to formalize.

What I mean to say is also that:

- ethical notions of unfairness should be contestable, since uncontestable notions of unfairness belong in the realm of ideology;
- legal notions of unfairness must be sufficiently demarcated to enable both foreseeability and contestability; and
- computational notions of unfairness must be formalizable, since one cannot train an algorithm without providing it with a machine-readable task and a performance metric.

Note that I have shifted from addressing fairness to addressing unfairness, because in a design context it may be a bit pretentious to claim that one can design ‘fairness’, whereas a sustained and systematic effort to design against unfairness will also keep us alert to new types of unfairness. Binary logic fails us here; the fact that something is not unfair (in some particular sense of the term) does not imply that it is fair (in all senses of the term). Fairness is what Gallie would term an essentially contested concept that requires vigilance and acuity rather than closure.

The point of this exercise is to develop mutual respect for the difference between ethical, legal, and computational notions of fairness and unfairness. To demonstrate what I mean with such mutual respect, I will sketch three approaches to the use of the COMPAS software: an ethical ‘by design’ approach, a legal ‘by design’ approach, and a computational ‘by design’ approach. Before doing so, I explain the background of the decisions supported by COMPAS.

11.3.2.1 The case of COMPAS

When deciding about whether to detain or release a criminal defendant or a criminal offender, courts in the United States assess the likelihood of recidivism. This may concern pre-trial decisions (probation), trial decisions (sentencing), and post-trial decisions on early release (parole). These decisions are to some extent discretionary, meaning the court is not bound by strict legal conditions (this may differ per state, and for sentencing stricter rules may apply). A high likelihood of recidivism is one of the factors weighing in on a decision to detain or release the defendant (who is awaiting trial), or of the offender (who was convicted and awaits sentencing or has been detained but is eligible for early release).

The idea is that detention prevents additional offences, so the goal of this particular assessment is to protect potential victims (this is often identified as protecting ‘the public’ or ‘the community’). In the case of a defendant the goal cannot be punishment, because being a *defendant* means there is no conviction yet. In the case of an *offender*, the goal of detention is punishment, early release can, for example, be justified as a reward for good behaviour, a way to reduce pressure on prisons, or a way to contribute to reintegration into society. These decisions, however, are not only based on the assessment of potential recidivism, they should also take into account what would be best for the defendant or offender.

On the website of the US Justice department,² the status and the goals of *parole* are clarified as follows:

When someone is paroled, they serve part of their sentence under the supervision of their community. The law says that the U.S. Parole Commission may grant parole if (a) the inmate has substantially observed the rules of the institution; (b) release would not depreciate the seriousness of the offense or promote disrespect for the law; and (c) release would not jeopardize the public welfare.

Parole has a three-fold purpose: (1) through the assistance of the United States Probation Officer, a parolee may obtain help with problems concerning employment, residence, finances, or other personal problems which often trouble a person trying to adjust to life upon release from prison; (2) parole protects society because it helps former prisoners get established in the community and thus prevents many situations in which they might commit a new offense; and (3) parole prevents needless imprisonment of those who are not likely to commit further crime and who meet the criteria for parole. While in the community, supervision will be oriented toward reintegrating the offender as a productive member of society.

Courts have been assessing the risk of recidivism based on hearing the defendant or offender and a whole range of further information is taken into account, not merely the recidivism likelihood. This seems to get lost in the discussion, and though this may be caused by the fact that the lofty wordings above reflect intention but not reality, it is crucial to remember that recidivism should not be the only criterion to decide on detain-or-release decisions.

² <https://www.justice.gov/uspc/frequently-asked-questions#q1>.

The assessment of the likelihood of recidivism is done by whoever is competent to decide on detention or release. Those competent (often courts, e.g. supported by parole boards, probation officers etc.) can use their common sense and their trained intuition as well as empirical reporting by experienced or expert advisers to reach a conclusion. In line with calls for ‘evidence based’ sentencing decisions, various types of data-driven software tools have been developed that are usually claimed to assess the relevant risk more accurately or more expediently. Some of this software has been developed by federal or state courts, but some courts rely on proprietary software from commercial vendors. One such vendor, with a substantial ‘market share’ was Northpointe (now Equivant), who developed the COMPAS system, which stands for correctional offender management profiling for alternative sanctions. The COMPAS risk score is based on six features, after its learner algorithm was trained on available data sets with a feature space of 137 features. The learner algorithm has found these six features highly indicative of recidivism. The risk score is based on an interview and/or a questionnaire that is filled in by the defendant or offender, and on their criminal file.

Because of the major impact of the use of proprietary software on detention decisions, Julia Angwin (an investigative journalist working with ProPublica), decided to test the accuracy of the predictions and came to the following conclusions (based on her own scientific data-driven research):

In forecasting who would re-offend, the algorithm correctly predicted recidivism for black and white defendants at roughly the same rate (59 percent for white defendants, and 63 percent for black defendants) but made mistakes in very different ways. It misclassifies the white and black defendants differently when examined over a two-year follow-up period.

Our analysis found that:

Black defendants were often predicted to be at a higher risk of recidivism than they actually were. Our analysis found that black defendants who did not recidivate over a two-year period were nearly twice as likely to be misclassified as higher risk compared to their white counterparts (45 percent vs. 23 percent).

White defendants were often predicted to be less risky than they were. Our analysis found that white defendants who re-offended within the next two years were mistakenly labeled low risk almost twice as often as black re-offenders (48 percent vs. 28 percent).

The analysis also showed that even when controlling for prior crimes, future recidivism, age, and gender, black defendants were 45 percent more likely to be assigned higher risk scores than white defendants.

Black defendants were also twice as likely as white defendants to be misclassified as being a higher risk of violent recidivism. And white violent recidivists were 63 percent more likely to have been misclassified as a low risk of violent recidivism, compared with black violent recidivists.

The violent recidivism analysis also showed that even when controlling for prior crimes, future recidivism, age, and gender, black defendants were 77 percent more likely to be assigned higher risk scores than white defendants.

This gave rise to a turbulent debate, where Northpointe accused Angwin of methodological incompetence, stating that their own predictions were the result of objective application of statistics. This in turn generated a series of scientific articles on both sides of the debate, resulting in a number of initiatives on the side of law, social science, and computer science to counter what has been termed ‘bias in machine learning’, finally prompting a new ACM conference dedicated to ‘fair accountable and transparent’ computing.

At some point, an offender was sentenced to six years of imprisonment, after the judge had taken note of the high risk score attributed by COMPAS.³ The offender, Eric Loomis, appealed the decision on the grounds that his sentence was based on proprietary software that should not have informed the decision because it was not possible to assess its accuracy, thereby violating his due process rights and/or because it may have wrongly taking gender into account. The appeals court rejected his appeal.

Note that the COMPAS recidivism risk score is part of the so-called ‘Presentence Investigation Report (PSR)’, that was used to determine the sentence. The PSR explicitly stated:

For purposes of Evidence Based Sentencing, actuarial assessment tools are especially relevant to: 1. Identify offenders who should be targeted for interventions. 2. Identify dynamic risk factors to target with conditions of supervision. 3. It is very important to remember that risk scores are not intended to determine the severity of the sentence or whether an offender is incarcerated (emphasis added by the court).

³ See the judgment of the Supreme Court of Wisconsin, 881 N.W.2d 749 (Wis. 2016), available at: <https://www.scotusblog.com/wp-content/uploads/2017/02/16-6387-op-bel-wis.pdf>.

The court of appeal, however stated:

In addition, the COMPAS report that was completed in this case does show the high risk and the high needs of the defendant. There's a high risk of violence, high risk of recidivism, high pre-trial risk; and so all of these are factors in determining the appropriate sentence.

(...)

You're identified, through the COMPAS assessment, as an individual who is at high risk to the community.

In terms of weighing the various factors, I'm ruling out probation because of the seriousness of the crime and because your history, your history on supervision, and the risk assessment tools that have been utilized, suggest that you're extremely high risk to re-offend.

The Supreme Court of Wisconsin did not overturn the decision of the court of appeal, stating that the high risk-score was corroborated by other evidence, basically concluding that the court would have made the same decision even if it had not seen the COMPAS assessment.

11.3.2.2 A computational 'fairness by design' approach to detain/release court decisions

There are three issues here:

1. the question whether the COMPAS output algorithm is indeed *accurate*, and what this means from a computational perspective;
2. the question whether the algorithm is *unfair*, and if so, what this means—in terms of computational formalization;
3. the question whether the answers to the previous questions are *objective*, and if so in what sense.

Julia Angwin's main point is that, though the accuracy for black and white defendants is the same, the error in the case of black defendants concerns *false positives* (they are given a higher risk-score compared to their actual recidivism), whereas in the case of white defendants the error concerns *false negatives* (they are given a lower risk-score compared to their actual recidivism). Northpointe/Equivant has argued that this is inevitable because black people (as an aggregate) recidivize more often. Proper use of

statistics—according to Northpointe/Equivant—results in an undesirable but unavoidable disparate outcome.

One could retort that this depends on how you train your learner algorithm. If the machine-readable task is to ensure that all defendants who do not recidivize will have the same error rate for both false positive and false negatives in the case of both black and white defendants, then the learner algorithm will learn just that.

There may be a ‘cost’ insofar as this may result in more false negatives for black people who do recidivize, but a ‘cost’ will actually be inevitable, it is inherent in the employment of statistics. The question of which cost we accept is not a matter of accuracy or objectivity, but of either ethics or law (and, obviously, the political choices made when writing the law).

This relates to the issue of fairness. Having concluded that statistics in itself does not dictate the machine learning research design choices made by Northpointe/Equivant, we suddenly find ourselves in the realm of fairness. Some may find it fair that a black person who will not recidivize has a higher chance of being detained due to a false positive than a white person, whereas they would find it unfair—or maybe dangerous—to make design choices that would result in a higher chance of false negatives for black people who will recidivize.

The underlying question is whether *it is unfair* to judge a black person based on the fact that other black people (according to the data) more often recidivize than white people, or whether *it is unfair* that a white person who will recidivize benefits from the fact that generally speaking (according to the data) white people recidivize less often than black people. In this case, we can’t have our cake and eat it too, a choice will have to be made between these two types of unfairness.

From a computer science perspective, both can be formalized and made operational. Due to the fact that as a society, we may not agree on the choice to be made here, it is difficult to demand ‘closure’ from computer scientists.

What they can do is:

- to explain the implications of the design choices and their trade-offs; and
- to develop still further and other ways to train a learner algorithm in ways that could reduce similar types of unfairness.

At this moment, computer scientists have come up with dozens of different ways to formalize fairness. This demonstrates that the employment of this type of software may seem expedient and effective, whereas in point of fact it may create more problems than it solves.

This conclusion may also be drawn from the Supreme Court of Wisconsin, where it finds that the appeal court would have made the same decision if COMPAS had not been used. Interestingly, computer science research by Farid and Dressel led them to the conclusion that the COMPAS algorithm does not outperform a randomly chosen set of human assessors who based their assessment on seven features. In other words, investing in this type of software may have no added value. Northpointe/Equivant, however, was seen to be rather proud that they did about as well as human assessors, arguing that their accuracy would improve (with more data). The Supreme Court of Wisconsin seems to assume the same, as it urged courts to adopt more evidence-based decision-support tools, though cautioning about the current state of the art.

11.3.2.3 An ethical ‘fairness by design’ approach to detain/release court decisions

When reading the research presented by Julia Angwin, Northpointe/Equivant, a number of other authors, and the *Loomis* case, one cannot but conclude that merely ‘fixing’ the COMPAS algorithms will not suffice.

During tutorials at different computer science conferences, Narayanan has presented *over twenty different formalizable definitions of fairness*, and in the bibliography below I refer to the draft version of a book he is co-authoring with Barocas and Hardt on *Fairness and Machine Learning. Limits and Opportunities*. Clearly, the more sophisticated the arguments of computer scientists for various types of fairness, the more we need to sit down and come to terms with the kind of fairness we should apply in what circumstances. This not only concerns the COMPAS software, but the employment of many other types of decision-support systems, such as predictive policing, taxation and social security fraud detection, eligibility for care (think of potentially abused children or the elderly), access to education, the job market, and insurance.

The case of COMPAS thus nicely demonstrates the complexity of the decisions that must be made by the court and of the interaction between different factors that play out on the side of the defendant or offender. In the

case of *Loomis*, the defendant had agreed to a plea bargain, which means that—even though he did not confess—he was willing to accept punishment. This is a common practice in the United States that offers the justice system some relief from procedural requirements, traded against a lowering of the sentence or fine for the defendant. The deal is struck between the public prosecutor and the defendant, meaning that the court is not bound by it, though most often taking it into account (some call this ‘trading with justice’). It may be that much of the unfairness starts here, and even much earlier, where black Americans have a much higher chance of being disadvantaged in numerous ways and of being treated in ways that do not reflect the idea that a government should treat each and every citizens with ‘equal concern and respect’.

Defining unfairness in a way that amends for both prejudice and for the result of previous unfair treatment and other root causes of recidivism is not an easy task, whether the assessment of the likelihood of recidivism is done by a human or a computational system. In both cases, the problem sits in the shift from an assessment at the aggregate level to the individual level (in psychology this is called stereotyping), and medical research tells us that what is reasonable at the level of epidemiology may be off-key at the individual level.

Let’s remind ourselves that we are making decisions like this, based on various types of generalization, every day. There is no way we can escape from the dilemma these decisions pose.

Here, I believe, the contribution of ethics can be pivotal. This will only work if we steer free from uninformed utilitarian cost-benefit analyses that weigh, for example, public goods such as privacy as if they are merely private interests, against private interests of the state under the heading of public security, often remaining stuck in simplistic act-utilitarianism. Similarly, we should not fall into the trap of romanticizing the singularity of individual defendants, claiming they should never be compared to others. As I have tried to elucidate in section 11.1, ethics is deeply concerned with the need to articulate rules that are not informed by parochial interests, both in rule-utilitarianism and in deontological reasoning. A naive interpretation of the rule that maximizes utility (aggregate or average) would possibly align with the position taken by Northpointe/Equivant, insofar as the cost of false positives of black defendants that would not recidivize were to be less than the cost of false negatives of black defendants that do recidivize. This position is naive

because the distribution of the cost is not taken into account (whose costs are weighted against whose benefits?), and also because this approach reinforces existing bias and may incur enormous cost down the line where black communities are confronted with a downward spiral of disrespect. Instead, we could investigate whether Rawls' maximin principle could be applied here, suggesting that fair algorithms should at least prevent loss of utility for the least advantaged, or develop a threshold in the learning algorithm that rules out picking on those already suffering systemic disadvantage.

But, maybe, the role of ethics is not only to achieve something like 'counter optimization'. Perhaps, virtue ethics and pragmatist ethics can highlight the need for human judgment, showing that in the end this may be less complicated and less dependent on invisible computation, while it can be called out in a more transparent way. Though the Wisconsin Supreme Court judged that due process was not violated, the mere fact that the problem can be articulated in terms of due process may help to frame the issue.

The court seems to give the COMPAS software the benefit of the doubt, hoping it will soon be better and admonishing courts in general to rely more rather than less on what it calls evidence-based sentencing. As one of the judges writes in her concurring opinion, however, the court allows the usage of these kinds of tools notwithstanding the observation that no agreement exists as to the reliability of COMPAS, neither in the scientific literature nor in the popular press. At some point, the tables may be turned, if current case law is overruled. Providing arguments based on ethical inquiry that takes into account the tension between individual retribution and equal treatment should help both legislatures and courts to refine their enactments and judgments, paying keen attention to the redistribution of disadvantages that may unintentionally occur due to disparate treatment.

11.3.2.4 A legal 'fairness by design' approach to detain/release court decisions

As indicated above (section 11.3.1), I believe that the legal obligation to incorporate DPbD *in the light of risks to the rights and freedoms of natural persons* is not restricted to privacy by design (and not even restricted to data subjects). On the contrary, the articulation in the GDPR emphasizes the need to foresee implications for other fundamental rights, as required by the DPIA.

This means that we already have a legal obligation to at least remedy 'unfairness by design'.

A court decision to detain or release a defendant or offender is most often *discretionary*; it is based on a broader margin of appreciation than other decisions, notably the conviction itself (due to the presumption of innocence, a court may not convict a person if there is reasonable doubt whether the defendant committed the offence). Under the rule of law, however, discretion is not equivalent with arbitrary decisionism. A court will have to consider a number of factors before coming down with a decision, and this consideration cannot be outsourced to a machine. The reason is that such outsourcing might on the one hand enable the scaling and the streamlining of decisions, but on the other hand it may deskill the judge to the extent that they are no longer required to actually consider these factors themselves, face to face with the defendant or offender. This may diminish the practical wisdom of the court, which increases the chance that courts will uncritically rely on the calculations of software they cannot assess.

This means that a ‘fairness by design’ approach in law requires two caveats:

1. To claim that an algorithm can ‘make’ decisions fair is overstating what algorithms can do in this space; for that reason, it is better to develop a ‘countering unfairness by design’ approach.
2. These tools should not be used to replace legal judgment but to challenge it, thus enhancing the practical wisdom of the court instead of diminishing it; for that reason, lawyers and computer scientists should sit down together to write code that keeps courts nimble and sharp.

11.4 Closure: The Force of Technology and the Force of Law

In this chapter, I have argued that if ethics aligns with the force of technology, the rule of law confronts a dangerous competitor in our normative space. The fact that ethics lacks the checks and balances of the rule of law signifies that we should not become overdetermined by ‘ethical technologies’ (whatever that could mean).

However, we can also imagine the use of technological affordances to limit the unfairness of algorithmic decision-making, thus underpinning the equal concern and respect that a government owes each and every one of its citizens. This will only work if algorithmic decision-support systems challenge the acuity of human judgment instead of replacing it.

References

On utilitarianism, deontological moral philosophy, virtue ethics, and pragmatism

- Alexander, Larry and Michael Moore. 'Deontological Ethics'. *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/win2016/entries/ethics-deontological/>.
- Hooker, Brad. 'Rule Consequentialism'. *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/win2016/entries/consequentialism-rule/>.
- Hursthouse, Rosalind and Glen Pettigrove. 'Virtue Ethics'. *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition), Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/win2018/entries/ethics-virtue/>.
- Johnson, Robert and Adam Cureton. 'Kant's Moral Philosophy'. *The Stanford Encyclopedia of Philosophy* (Spring 2019 Edition), Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/spr2019/entries/kant-moral/>.
- Legg, Catherine and Christopher Hookway. 'Pragmatism'. *The Stanford Encyclopedia of Philosophy* (Spring 2019 Edition), Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/spr2019/entries/pragmatism/>.
- Sinnott-Armstrong, Walter. 'Consequentialism'. *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/win2015/entries/consequentialism/>.
- Varela, Francisco J. 1992. *Ethical Know-How*. Stanford: Stanford University Press.

On ethics in design

- Awad, Edmond, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. 2018. 'The Moral Machine Experiment'. *Nature* 563 (7729): 59. <https://doi.org/10.1038/s41586-018-0637-6>.
- Dignum, Virginia. 2018. 'Ethics in Artificial Intelligence: Introduction to the Special Issue'. *Ethics and Information Technology* 20 (1): 1–3. <https://doi.org/10.1007/s10676-018-9450-z>.
- Hoven, Jeroen van den, Pieter E. Vermaas, and Ibo van de Poel, eds. 2015. *Handbook of Ethics, Values, and Technological Design: Sources, Theory, Values and Application Domains*. 2015 ed. Dordrecht: Springer.
- Nissenbaum, Helen Fay. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford: Stanford Law Books.
- Porcaro, Keith. 2019. 'Detain/Release: Simulating Algorithmic Risk Assessments at Pretrial'. *Medium* (blog). 8 January 2019. <https://medium.com/berkman-klein-center/detain-release-simulating-algorithmic-risk-assessments-at-pretrial-375270657819>.

- Powles, Julia. 2018. 'The Seductive Diversion of "Solving" Bias in Artificial Intelligence'. *Medium*. 7 December 2018. <https://medium.com/s/story/the-seductive-diversion-of-solving-bias-in-artificial-intelligence-890df5e5ef53>.
- Wagner, Ben. 2018. 'Ethics as an Escape from Regulation. From "Ethics-Washing" to "Ethics-Shopping?"' In *Being Profiled: Cogitas Ergo Sum: 10 Years of Profiling the European Citizen*, edited by Emre Bayamlioglu, Irina Baraliuc, Lisa Janssens, and Mireille Hildebrandt, 84–87. Amsterdam: Amsterdam University Press.

On fair computing and framing problems

- Barocas, Solon, and Andrew D. Selbst. 2016. 'Big Data's Disparate Impact'. *California Law Review* 104: 671–732.
- Barocas, Solon, Moritz Hardt, and Arvind Narayanan, draft version of *Fairness and Machine Learning. Limitations and Opportunities*. <https://fairmlbook.org/pdf/fairmlbook.pdf>.
- Callon, M., and J. Law. 2005. 'On Qualculation, Agency, and Otherness'. *Environment and Planning D: Society and Space* 23 (5): 717–33.
- Chouldechova, Alexandra. 2017. 'Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments'. *Big Data* 5 (2): 153–63. <https://doi.org/10.1089/big.2016.0047>.
- Chouldechova, Alexandra, and Aaron Roth. 2018. 'The Frontiers of Fairness in Machine Learning'. *ArXiv:1810.08810 [Cs, Stat]*, October. <http://arxiv.org/abs/1810.08810>.
- Dressel, Julia, and Hany Farid. 2018. 'The Accuracy, Fairness, and Limits of Predicting Recidivism'. *Science Advances* 4 (1): eaao5580. <https://doi.org/10.1126/sciadv.aao5580>.
- Equivant. 2018. 'Response to ProPublica: Demonstrating Accuracy Equity and Predictive Parity'. *Equivant* (blog). 1 December 2018. <https://www.equivant.com/response-to-propublica-demonstrating-accuracy-equity-and-predictive-parity/>.
- Equivant. 2018. 'Official Response to Science Advances'. *Equivant* (blog). 18 January 2018. <https://www.equivant.com/official-response-to-science-advances/>.
- Fricker, Miranda. 2007. 'Hermeneutical Injustice'. In *Epistemic Injustice: Power and the Ethics of Knowing*, 147–75. Oxford: Oxford University Press. <https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780198237907.001.0001/acprof-9780198237907-chapter-8>.
- Gallie, W.B. 1956. 'Essentially Contested Concepts'. *Proc. Aristotelian Soc'ty* 56: 167–98.
- Kroll, Joshua, Joanna Huey, Solon Barocas, Edward Felten, Joel Reidenberg, David Robinson, and Harlan Yu. 2017. 'Accountable Algorithms'. *University of Pennsylvania Law Review* 165 (3): 633.

Northpointe. 2012. *Practitioners Guide to COMPAS*. http://www.northpointeinc.com/files/technical_documents/FieldGuide2_081412.pdf.

On the inner morality of the Rule of Law (and Rule of Law in cyberspace)

Brownsword, Roger. 2016. 'Technological Management and the Rule of Law'. *Law, Innovation and Technology* 8 (1): 100–40. <https://doi.org/10.1080/17579961.2016.1161891>.

Dworkin, Ronald. 1991. *Law's Empire*. Glasgow: Fontana.

Hildebrandt, Mireille. 2015. 'Radbruch's Rechtsstaat and Schmitt's Legal Order: Legalism, Legality, and the Institution of Law'. *Critical Analysis of Law* 2 (1). <http://cal.library.utoronto.ca/index.php/cal/article/view/22514>.

Rawls, John. 2005. *A Theory of Justice*. Cambridge, MA: Belknap Press.

Reed, Chris, and Andrew Murray. 2018. *Rethinking the Jurisprudence of Cyberspace*. Cheltenham: Edward Elgar.

Waldron, Jeremy. 2011. 'The Rule of Law and the Importance'. *Nomos* 50: 3–31.