AN INTERDISCIPLINARY TEXTBOOK

Introduction to

Digital Economics and Entrepreneurship

HENRIETTA NAGY JAROSŁAW KINAL TAMÁS TÓTH

Henrietta Nagy
Jarosław Kinal
Tamás Tóth

INTRODUCTION TO DIGITAL ECONOMICS AND ENTREPRENEURSHIP

An interdisciplinary textbook for students and researchers
in economics, sociology and media studies

INTRODUCTION TO DIGITAL ECONOMICS AND ENTREPRENEURSHIP

An interdisciplinary textbook for students and researchers in economics, sociology and media studies

Autors:

Henrietta Nagy, Jarosław Kinal, Tamás Tóth

Reviewer:

György Neszmélyi

Publisher:

MILTON FRIEDMAN EGYETEM

Milton Friedman University, 1039 Budapest, Kelta utca 2, Hungary

ISBN 978-963-9559-80-6

Budapest 2026
First Edition

This is an Open Access publication distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0) (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and the source are properly credited.

Table of Contents

Preface 8
How to Use This Book 9
List of Abbreviations 10
Introduction: Why Digital Economics Needs an Interdisciplinary Lens 12
Chapter 1. The Digital Economy: Concepts, Foundations, and Historical Trajectories 16
1.1. Defining the Digital Economy: A Contested Concept 16
1.2. Historical Trajectories: Five Waves of Digitalization 18
1.3. The Economic Properties of Information and Digital Goods 19
1.4. Network Effects, Increasing Returns, and Platform Logic 21
1.5. Measuring the Digital Economy and the Productivity Paradox 22
1.6. Digital Entrepreneurship: A First Approximation 23
1.7. Three Disciplinary Lenses: Economics, Sociology, Media Studies 24
1.8. Chapter Summary and Discussion Questions 26
1.9. Case Studies and Debates 28
1.10. In Depth: The Political Arithmetic of Measuring Digitalization 31
1.11. In Depth: A Working Chronology of the Digital Economy, 1945-2025 34
Chapter 2. The Economics of Information and Digital Goods 37
2.1. Information as an Economic Good: Non-Rivalry and Excludability 37
2.2. Zero Marginal Cost, Versioning, and Price Discrimination 39
2.3. The Five Cost Reductions: Search, Replication, Transportation, Tracking, Verification 40
2.4. Intellectual Property and Open Alternatives 42
2.5. Chapter Summary and Discussion Questions 43
2.6. Case Studies and Debates 45
2.7. In Depth: Search, Matching, and the Economics of Recommendation 48
2.8. In Depth: Disclosure, Certification, and the Market for Quality 51
Chapter 3. Platforms, Network Effects, and Two-Sided Markets 53
3.1. Direct and Indirect Network Effects 53
3.2. Two-Sided and Multi-Sided Markets 54
3.3. Winner-Take-Most Dynamics, Lock-In, and Switching Costs 56
3.4. Platform Envelopment and Ecosystem Competition 58
3.5. Chapter Summary and Discussion Questions 59
3.6. Case Studies and Debates 61
3.7. In Depth: The Platform as Regulator - Governance as Institutional Design 65
3.8. In Depth: The Economics of Interoperability 67
Chapter 4. Digital Entrepreneurship: Theory, Actors, and Ecosystems 69
4.1. From Schumpeter to Digital Affordances 69
4.2. Lean Startups, Scalability, and Born-Global Ventures 71

Applied Toolkit: From Opportunity to Product-Market Fit 73
4.3. Entrepreneurial Ecosystems and Venture Finance 74
4.4. Social, Cultural, and Institutional Conditions of Digital Venturing 75
4.5. Chapter Summary and Discussion Questions 77
4.6. Case Studies and Debates 80
4.7. In Depth: Venture Finance Through the Cycle 83
4.8. In Depth: Cooperative and Social Alternatives - Owning the Platform Otherwise 86
Chapter 5. Business Models of the Digital Economy 88
5.1. Advertising-Funded, Subscription, and Freemium Models 88
5.2. Marketplaces, Aggregators, and On-Demand Services 90
5.3. Software-as-a-Service and Cloud Infrastructures 91
5.4. The Creator Economy and Direct Monetization 92
Applied Toolkit: Reading a Digital Business Model 93
5.5. Chapter Summary and Discussion Questions 94
5.6. Case Studies and Debates 96
5.7. In Depth: Pricing the Free - The Machinery of Digital Advertising 100
5.8. In Depth: Play Money - The Economics of Games and Virtual Goods 103
5.9. In Depth: Money as a Platform - Payments, Open Banking and Digital Currencies 105
Chapter 6. Data, Algorithms, and Artificial Intelligence as Economic Resources 106
6.1. Data as Capital: Extraction, Aggregation, and Valuation 107
6.2. Algorithmic Decision-Making and Prediction Machines 108
6.3. Surveillance Capitalism and Its Critics 109
6.4. Generative AI and the New Automation Debate 110
Applied Extension: The Economics of Generative AI and AI-Native Firms 112
6.5. Chapter Summary and Discussion Questions 114
6.6. Case Studies and Debates 116
6.7. In Depth: The Production Economics of Machine Learning 119
6.8. In Depth: The Economics of Privacy 121
Chapter 7. Digital Labour and the Transformation of Work 124
7.1. Gig Work, Crowdwork, and Algorithmic Management 124
7.2. Skills, Polarization, and Technological Unemployment 126
7.3. Invisible Labour: Moderation, Microwork, and Free Labour 127
7.4. Chapter Summary and Discussion Questions 129
7.5. Case Studies and Debates 131
7.6. In Depth: Remote Work and the New Geography of Labour 135
7.7. In Depth: Regulating Platform Work - Europe’s Presumption of Employment 137
Chapter 8. The Sociology of Digital Markets: Trust, Networks, and Inequality 138
8.1. Embeddedness and Trust in Online Exchange 139
8.2. Digital Divides and Stratified Participation 140

8.3. Reputation Systems and the Quantification of Social Life 141
8.4. Chapter Summary and Discussion Questions 143
8.5. Case Studies and Debates 145
8.6. In Depth: Designing Reputation - A Mechanism Designer's Notebook 148
8.7. In Depth: From Persons to Systems - The Sociology of Institutional Trust 150
Chapter 9. Media, Attention, and the Political Economy of Communication 152
9.1. The Attention Economy and Advertising Markets 152
9.2. Platformization of Media Industries and Journalism 154
9.3. Algorithmic Curation, Public Spheres, and Misinformation 155
9.4. Chapter Summary and Discussion Questions 157
9.5. Case Studies and Debates 160
9.6. In Depth: Metrics at Work - How Measurement Remade the Newsroom 163
9.7. In Depth: The Pipes - A Political Economy of Internet Infrastructure 165
Chapter 10. Regulation, Governance, and the Future of the Digital Economy 167
10.1. Competition Policy for Digital Markets 167
10.2. Data Protection, Platform Regulation, and the European Model 169
10.3. Global Governance, Digital Sovereignty, and Development 170
10.4. Scenarios for the Next Decade 172
10.5. Chapter Summary and Discussion Questions 173
10.6. Case Studies and Debates 176
10.7. In Depth: Governing Intelligence - Comparative Architectures of AI Regulation 180
Conclusion: The Digital Economy as a Constructed Order 183
Epilogue: Reading the Next Wave 186
Capstone Project: Audit a Digital Market 187
Glossary of Core Concepts 188
Subject Guide: Core Topics Across the Book 190
References 191

Preface

This book began as three sets of lecture notes that refused to stay apart. Teaching the economics of digital markets to economics students, the sociology of platforms to sociology students, and the political economy of media to students of communication, the authors found the same firms, the same mechanisms, and the same controversies appearing in all three courses - described in three vocabularies that barely acknowledged one another. The economics students could derive the two-sided price structure but had no language for what the zero price does to the public sphere; the sociology students commanded the critique of platform power but treated its economic mechanics as a black box; the media students knew the attention economy as lived experience and normative problem but not as auction design. Each cohort, in short, held a piece of an object that exists only whole. The book is the attempt to teach the whole object.

Its intellectual debts are recorded in the references, but three should be named at the outset because they supplied the book's architecture rather than its details: the economics of information of Shapiro and Varian, which demonstrated that the digital economy's strangeness is derivable from a small set of durable principles; the embeddedness tradition of Granovetter, which insists that markets are social institutions all the way down; and the political economy of communication running from Smythe to the contemporary platform-studies literature, which never allowed the analysis to forget that the infrastructures of exchange are also the infrastructures of public life. Where the book has a method, it is the disciplined alternation among these three: mechanism, institution, medium.

A textbook on a moving object incurs obligations of candour. The empirical claims in these pages - market structures, regulatory states of play, technological frontiers - were accurate to the best of the authors' knowledge at the time of writing and will decay at the digital economy's usual pace; readers should treat them as dated snapshots and check the current state before relying on them. The conceptual claims are built to last longer, and the reader who finds a case study superseded is invited to treat the supersession itself as an exercise: the book's mechanisms should explain what changed. All references follow the seventh edition of the Publication Manual of the American Psychological Association; the authors welcome corrections. Responsibility for the errors that remain is, in the traditional and accurate formula, the authors' alone.

How to Use This Book

The book is designed for a semester's teaching but built for three different semesters. Its ten chapters divide, roughly, into an economic core (Chapters 1-3), an entrepreneurship and business-model sequence (Chapters 4-5), a data and labour sequence (Chapters 6-7), and a sociological and media sequence (Chapters 8-9), with regulation (Chapter 10) as the convergence point. Chapter 1 is the common foundation and should open any course based on the book; Chapter 10 presupposes the whole and should close it. Between these, the dependencies are few and explicit: Chapter 3 requires Chapter 2; Chapter 5 requires Chapter 3; Chapters 6 through 9 each require Chapters 1 through 3 but may be taken in any order; Chapter 4 may follow Chapter 1 directly in courses that foreground entrepreneurship.

Three course pathways have been tested in teaching. An *economics-led* course runs the chapters in printed order, treating Chapters 8 and 9 as the extension of price theory into social and communicative domains. A *sociology-led* course opens with Chapters 1, 7, and 8, then works backward through the mechanisms (Chapters 2 and 3) that structure the phenomena, closing with Chapters 6 and 10. A *media-and-communication* course opens with Chapters 1 and 9, uses Chapters 2, 3, and 5 to supply the economic machinery of the attention economy, and closes with Chapters 6 and 10. In mixed-audience courses - the book's native habitat - the printed order works, provided the instructor enforces the discipline the book preaches: economists present the sociological readings, and vice versa.

Each chapter carries a fixed pedagogical apparatus. Learning objectives open the chapter and are written to be assessable: examinations can be constructed directly against them. The chapter summary and discussion questions consolidate the argument; the discussion questions are calibrated for seminar preparation and short written assignments. The case studies and debates sections are the book's laboratory: cases apply the chapter's concepts to documented episodes and end with questions of increasing difficulty; debates state a genuinely contested motion with the strongest case on each side, and are designed to be staged - assign sides in advance, require the operative test or policy to be drafted in writing, and grade the quality of concession as well as of advocacy. A glossary of core concepts and a subject guide mapping topics across chapters support revision. Instructors assessing by coursework will find that each chapter's case studies support a two-thousand-word analytical essay, and that the debates convert directly into structured oral examinations.

PEDAGOGICAL ARCHITECTURE

Key Concepts. A compact map of the chapter's indispensable terms before the detailed argument begins.

Common Misconception. A deliberately chosen error or oversimplification that students should learn to detect.

Worked Example. A short numerical or decision problem that converts abstract mechanisms into an auditable calculation.

Three Lenses. The same phenomenon viewed through economics, sociology, and media/communication studies, preserving the book's mechanism-institution-medium architecture.

2026 Update. Used only where fast-moving evidence or regulation materially changes the interpretation of the chapter.

List of Abbreviations

AI	artificial intelligence
API	application programming interface
ARPU	average revenue per user
ARR	annual recurring revenue
CAC	customer acquisition cost
CBDC	central bank digital currency
CDN	content-delivery network
DMA	Digital Markets Act (Regulation (EU) 2022/1925)
DSA	Digital Services Act (Regulation (EU) 2022/2065)
EU	European Union
GDP	gross domestic product
GDPR	General Data Protection Regulation (Regulation (EU) 2016/679)
GPAI	general-purpose artificial intelligence (model)
GPT	general purpose technology (Ch. 1); also, in industry usage, generative pre-trained transformer
GSP	generalized second-price (auction)
ICT	information and communication technology
IEA	International Energy Agency
ILO	International Labour Organization
IP	intellectual property; also Internet Protocol
LTV	customer lifetime value

MiCA	Markets in Crypto-Assets Regulation (Regulation (EU) 2023/1114)
ML	machine learning
MRR	monthly recurring revenue
MVP	minimum viable product
NIS2	Directive (EU) 2022/2555 on a high common level of cybersecurity across the Union
OECD	Organisation for Economic Co-operation and Development
PSD2	second Payment Services Directive (Directive (EU) 2015/2366)
SaaS	software-as-a-service
SAM	serviceable available market
SOM	serviceable obtainable market
TAM	total addressable market
UBI	universal basic income
UNCTAD	United Nations Conference on Trade and Development
UOKiK	Office of Competition and Consumer Protection (Poland)
UPI	Unified Payments Interface (India)

Introduction: Why Digital Economics Needs an Interdisciplinary Lens

Few transformations in modern economic history have been as rapid, as pervasive, and as conceptually disorienting as the digitalization of economic life. Within the span of a single generation, digital technologies have moved from the periphery of the economy - a specialized sector producing computers and software - to its very core, reshaping how goods are produced and priced, how firms are founded and financed, how labour is organized and remunerated, and how markets, media, and publics interact. Estimates of the size of the digital economy vary considerably depending on definition, but even conservative measurements suggest that the digitally enabled share of value added in advanced economies has grown steadily for three decades, while the broadest conceptions - encompassing all economic activity that relies significantly on digital inputs - approach the economy as a whole (Bukht & Heeks, 2017; OECD, 2020). To study the digital economy today is thus increasingly to study the economy as such.

This book offers a systematic introduction to digital economics and digital entrepreneurship, written deliberately for three audiences that too rarely read the same texts: economists, sociologists, and media and communication scholars. The rationale for this triple address is not a rhetorical gesture toward interdisciplinarity but a substantive claim about the object of study itself. The digital economy is simultaneously an economic phenomenon - characterized by distinctive cost structures, network effects, and market dynamics; a social phenomenon - embedded in networks of trust, stratified by inequalities of access and skill, and productive of new forms of work and social organization; and a media phenomenon - for the dominant firms of the digital age are, in large measure, media companies whose revenues derive from the capture and resale of human attention, and whose infrastructures now condition public communication itself (Castells, 2010; van Dijck et al., 2018; Zuboff, 2019). No single discipline possesses the conceptual vocabulary adequate to this object.

Consider a deceptively simple question: what kind of firm is a social media platform? For the economist, it is a multi-sided market that matches advertisers with users and exhibits strong indirect network effects, raising classic questions about pricing, entry barriers, and competition policy (Rochet & Tirole, 2003). For the sociologist, it is an organization that restructures social ties, mediates the presentation of self, and produces new hierarchies of visibility and reputation. For the media scholar, it is a publisher-like intermediary that

algorithmically curates the flow of public information while disclaiming editorial responsibility, thereby transforming the political economy of journalism and the structure of the public sphere (van Dijck et al., 2018). Each answer is correct; each is incomplete. The premise of this book is that students and researchers who can move among these three registers will understand the digital economy better than those confined to any one of them.

The book pursues three aims. The first is conceptual: to equip readers with a rigorous vocabulary for analysing digital markets and digital ventures - non-rivalry, marginal cost, network externalities, two-sidedness, platform envelopment, algorithmic management, datafication, platformization, and related notions. Wherever possible, concepts are traced to their origins in the scholarly literature, so that readers encounter not only definitions but the intellectual debates from which those definitions emerged. The second aim is empirical: to acquaint readers with the principal findings of research on the digital economy, including the measurement debates, the productivity paradox, the economics of platforms, the transformation of work, and the restructuring of media industries. The third aim is critical: to present, alongside the celebratory accounts of digital innovation, the substantial body of scholarship that interrogates the concentration of market power, the extraction of behavioural data, the precarization of digital labour, and the consequences of algorithmic curation for democratic publics (Srnicek, 2017; Zuboff, 2019).

The treatment of entrepreneurship in this volume follows the same interdisciplinary logic. Digital entrepreneurship is examined not merely as a managerial practice - the founding and scaling of technology ventures - but as a socio-economic phenomenon with its own theory, drawing on Schumpeter's (1934, 1942) classic account of the entrepreneur as the agent of innovation and, in his later work, of creative destruction. It also draws on contemporary reformulations that emphasize the distinctive affordances of digital technologies: the fluidity of digital artefacts, the scalability afforded by cloud infrastructures, the generativity of open platforms, and the compression of the distance between idea and market (Nambisan, 2017; Autio et al., 2018). At the same time, the book attends to the institutional, cultural, and financial conditions under which digital ventures emerge, and to the geographically uneven distribution of entrepreneurial ecosystems.

As a textbook, this volume is designed for advanced undergraduate and graduate courses in economics, sociology, media studies, management, and interdisciplinary programmes in digital society. Each chapter opens with learning objectives, develops its argument through the relevant scholarly literature, and closes with a summary and discussion questions intended for seminar use. The bibliography privileges peer-reviewed journal articles, monographs from

academic presses, and reports of international organizations such as the OECD, the World Bank, and the European Commission; journalistic and industry sources are used sparingly and only where scholarly treatments are unavailable, as is inevitable in a field whose object changes faster than its literature.

A note on intellectual humility is warranted. The digital economy is a moving target. Several of the empirical claims reported in this book - market shares, adoption rates, regulatory frameworks - will age quickly, and readers are encouraged to treat them as snapshots rather than settled facts. The conceptual apparatus, by contrast, has proven more durable: the economics of information articulated by Shapiro and Varian (1999) at the close of the twentieth century remains a remarkably reliable guide to phenomena its authors could not have foreseen, and the sociological analysis of the network society developed by Castells (2010) continues to frame contemporary debates. It is on this conceptual durability that an introduction such as this one must rest.

The book unfolds in ten chapters. Chapter 1 establishes the foundations: it surveys competing definitions of the digital economy, reconstructs the historical waves of digitalization, introduces the distinctive economic properties of digital goods, and presents the three disciplinary lenses that organize the volume. Chapters 2 and 3 develop the economic core, treating the economics of information and the theory of platforms and network effects respectively. Chapters 4 and 5 turn to entrepreneurship: the theory and ecology of digital venturing, and the business models through which digital ventures capture value. Chapters 6 and 7 address the newest strata of the digital economy - data, algorithms, and artificial intelligence as economic resources, and the transformation of labour they entail. Chapters 8 and 9 foreground the sociological and media-theoretical perspectives, examining trust, inequality, and reputation in digital markets, and the attention economy and platformization of media industries. Chapter 10 concludes with regulation, governance, and scenarios for the coming decade.

Readers may traverse the book in several ways. Economists will find the most familiar terrain in Chapters 1 through 5 and may use Chapters 8 and 9 to extend their analysis into social and communicative domains. Sociologists may begin with Chapters 1, 7, and 8 before working backward into the economic mechanisms that structure the phenomena they study. Media and communication scholars may begin with Chapters 1 and 9, where they will recognize their concerns most directly, but are urged not to skip the economic foundations, for the political economy of platforms is unintelligible without the theory of two-sided markets. It is precisely

in these crossings - the sociologist reading price theory, the economist reading platform studies - that the book hopes to make its contribution.

Chapter 1. The Digital Economy: Concepts, Foundations, and Historical Trajectories

Learning objectives. After studying this chapter, the reader should be able to: (1) distinguish narrow, broad, and holistic definitions of the digital economy and explain why definitional choices matter for measurement and policy; (2) reconstruct the principal historical waves of digitalization from mainframe computing to the platform and artificial intelligence era; (3) explain the economic properties of information and digital goods, including non-rivalry, near-zero marginal cost of reproduction, and their consequences for pricing and market structure; (4) define network effects and increasing returns and relate them to the emergence of platform-dominated markets; (5) summarize the main approaches to measuring the digital economy and the debate known as the productivity paradox; (6) provide a working definition of digital entrepreneurship; and (7) articulate what economics, sociology, and media studies each contribute to the analysis of digital capitalism.

KEY CONCEPTS

digital economy; digitization, digitalization, and digital transformation; general-purpose technology; complementarity; productivity paradox; digital entrepreneurship

1.1. Defining the Digital Economy: A Contested Concept

The term digital economy entered wide circulation in the mid-1990s, popularized above all by Don Tapscott's (1996) book of that title, which announced an "age of networked intelligence" in which the digitization of information would transform commerce, work, and organization. Tapscott wrote at the moment when the World Wide Web was opening the Internet to commercial use, and his account - like Nicholas Negroponte's (1995) contemporaneous *Being Digital*, with its celebrated contrast between moving atoms and moving bits - belongs to a first wave of literature that treated the digital economy primarily as the economy of the Internet: electronic commerce, online services, and the firms that provided them. This narrow conception had the virtue of tractability. One could, at least in principle, count the firms, jobs, and revenues attributable to Internet-based activity and track their growth.

Three decades later, the narrow conception has become untenable, for digital technologies now permeate sectors - agriculture, logistics, banking, health care, public administration - that no one would classify as part of an "Internet industry." Reviewing dozens of definitions advanced by scholars, governments, and international organizations, Bukht and

Heeks (2017) proposed an influential three-tier framework that has structured subsequent debate. The *core* digital sector comprises the production of digital goods and services themselves: hardware manufacture, software and information services, telecommunications. The *digital economy* proper adds activities whose existence depends on digital technologies - digital platforms, the app economy, the sharing economy, and digitally delivered services. The *digitalized economy*, finally, encompasses all economic activity significantly enhanced by digital inputs, from precision agriculture to algorithmic trading - a scope that shades into the economy as a whole. On Bukht and Heeks's (2017) estimates, the digital economy in its narrow scope - the core sector together with digital and platform services - accounted for roughly five per cent of global GDP at the time of writing, while the broadest tier defied precise quantification altogether.

Definitional choices are not innocent. As the OECD (2020) has emphasized in its work toward a common measurement framework, the estimated size of the digital economy can vary by an order of magnitude depending on the tier adopted, with immediate consequences for taxation, competition policy, and development strategy. A government persuaded that the digital economy is a small, discrete sector will regulate it as it regulates any industry; a government persuaded that digitalization is a general transformation of the economy will treat digital policy as economic policy tout court. Economists have increasingly favoured the latter view, conceptualizing digital technology as a general purpose technology in the sense of Bresnahan and Trajtenberg (1995): a technology, like the steam engine or electricity, whose pervasiveness, capacity for improvement, and complementarity with downstream innovation allow it to reshape the entire productive apparatus over decades. Goldfarb and Tucker (2019), in their landmark review of digital economics, accordingly define the field not by a sector but by a question: what becomes different about economic activity when digital technology reduces the costs of search, replication, transportation, tracking, and verification?

For the purposes of this book, we adopt a working definition in the spirit of this literature: *the digital economy is the ensemble of economic activities, institutions, and relations in which digitized data and Internet-based connectivity function as essential inputs to the creation, distribution, and capture of value*. The definition is deliberately broad. It encompasses the core digital sector and the platform economy, but it also directs attention to the digitalization of incumbent industries and - crucially for the sociologist and the media scholar - to the institutions and relations, not merely the transactions, that digitalization reorganizes.

COMMON MISCONCEPTION

Digitalization is not synonymous with "the technology sector." A narrow ICT-sector definition measures producers of digital technologies; a broader economic definition follows digital inputs, infrastructures, data, and digitally mediated transactions across conventional sectors. The denominator therefore depends on the analytical question.

1.2. Historical Trajectories: Five Waves of Digitalization

The digital economy did not arrive at once; it accumulated in waves, each layering new capabilities, business models, and social arrangements upon its predecessors. A first wave, from the 1950s to the 1970s, was the era of the mainframe: computing as a scarce, centralized resource deployed by large corporations, governments, and universities for calculation and record-keeping. Its economic significance lay less in new products than in the rationalization of back-office processes - payroll, inventory, reservations - and it gave rise to the first great information technology firms and to the professional stratum of programmers and systems analysts whose emergence sociologists of the period were quick to register in debates over post-industrial society (Bell, 1973).

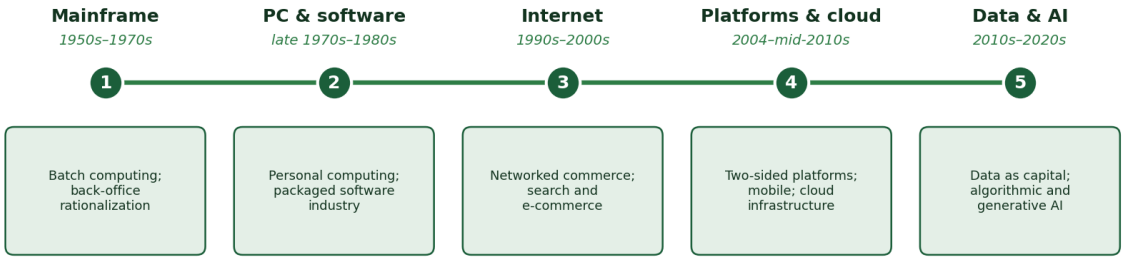

Figure 1.1. Five cumulative waves of digitalization and their dominant economic form.

The second wave, in the late 1970s and 1980s, brought the personal computer: computing as a decentralized tool on the desks of firms and, increasingly, households. The commodification of hardware, the rise of the packaged-software industry, and the emergence of the software entrepreneur as a cultural figure date from this period. The third wave, from the early 1990s, was the Internet wave proper. The opening of the network to commercial traffic and the invention of the World Wide Web created, for the first time, a universal, low-cost channel for the distribution of information goods and the coordination of exchange, unleashing the dot-com boom whose speculative excesses ended in the crash of 2000-2001. Carlota Perez's (2002) analysis of technological revolutions and financial capital reads that crash not as a

refutation of the digital economy but as a recurrent phase in the installation of a new techno-economic paradigm, in which financial euphoria overshoots the technology's short-run possibilities while nonetheless funding the infrastructure - in this case, fibre-optic networks and a generation of technical skill - on which the subsequent deployment period would draw.

The fourth wave, from roughly 2004 to the mid-2010s, is the age of the platform: Web 2.0 and user-generated content, social media, smartphones, cloud computing, and the multi-sided business models that converted connectivity into durable market power. It is in this period that the firms now dominating stock-market capitalizations attained their scale, and that scholars began to speak of platform capitalism (Srnicek, 2017) and the platform society (van Dijck et al., 2018) to signal that platforms had become not merely businesses but infrastructures of social life. The fifth wave, gathering force through the 2010s and accelerating dramatically in the 2020s, is the wave of data-driven artificial intelligence, in which machine learning systems trained on the data exhaust of the platform era are applied to prediction, recommendation, content generation, and automation across the economy (Brynjolfsson & McAfee, 2014; Agrawal et al., 2018). Each wave, it bears emphasis, presupposed its predecessors: there is no AI wave without the data of the platform wave, no platform wave without the connectivity of the Internet wave, and so on. The digital economy is best understood not as an event but as a cumulative, path-dependent process.

1.3. The Economic Properties of Information and Digital Goods

Why should digital markets behave differently from markets for wheat, steel, or haircuts? The canonical answer was given by Shapiro and Varian (1999) in *Information Rules*, still the indispensable primer on the economics of information: information goods are costly to produce but cheap to reproduce. The first copy of a film, a software application, or a scientific database may cost millions to create; each additional copy costs, in the digital case, effectively nothing. This cost structure - high fixed costs, near-zero marginal costs - has far-reaching consequences. Competitive pricing at marginal cost, the benchmark of textbook economics, would drive the price of information goods toward zero and make it impossible to recover the costs of creation. Producers of digital goods therefore price according to value rather than cost, engage in extensive versioning and price discrimination, bundle goods together, and seek institutional protections - intellectual property, technical protection measures - that restore a measure of excludability to goods that are naturally non-excludable.

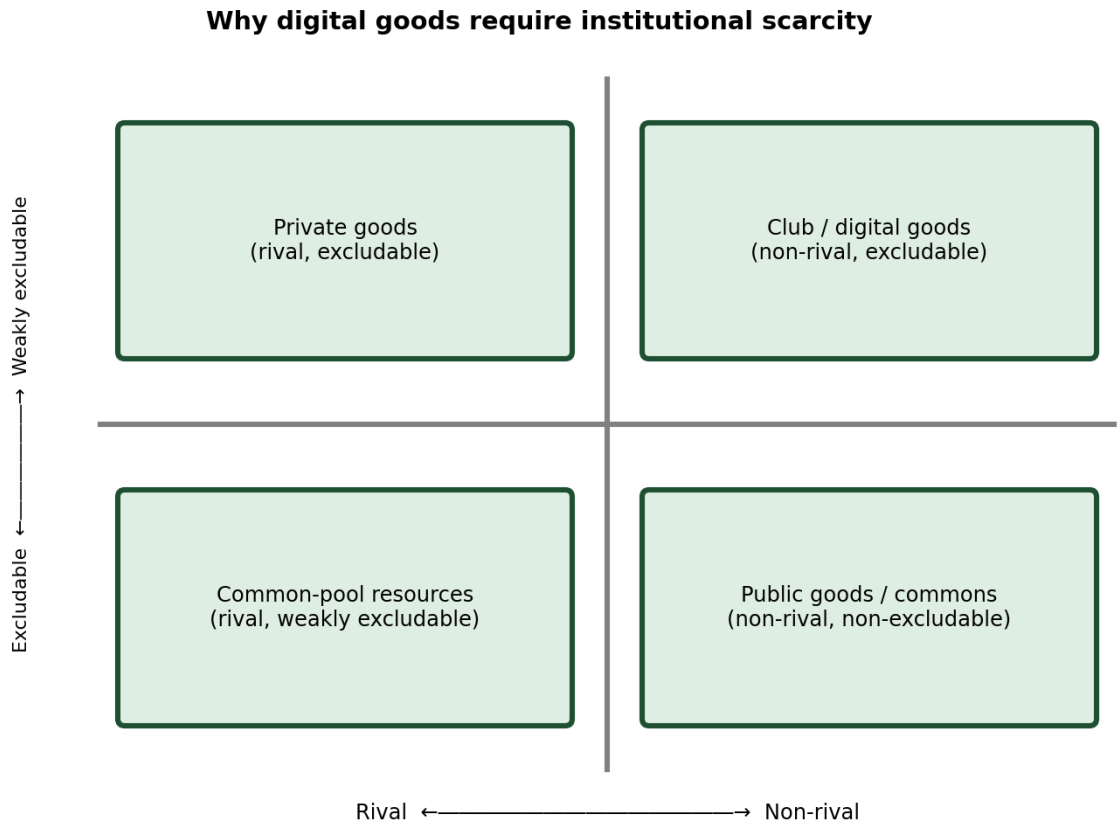

Figure 1.2. Rivalry and excludability: locating digital goods within the economic taxonomy of goods.

Two properties of information deserve particular emphasis. First, information is *non-rival*: one person's consumption of a digital file does not diminish its availability to others, in sharp contrast to rival goods such as food or machinery. Non-rivalry places digital goods in the neighbourhood of public goods and underlies both the chronic anxiety of content industries over piracy and the extraordinary phenomena of commons-based peer production - free and open-source software, Wikipedia - analysed by Benkler (2006), in which non-market cooperation among distributed volunteers produces information goods of the first economic importance. Second, information is an *experience good*: its value cannot be fully known until it has been consumed, at which point the consumer no longer needs to purchase it. Markets for information goods therefore rely heavily on branding, reviews, reputation systems, sampling, and free tiers - institutional devices for solving the problem of quality uncertainty that Akerlof (1970) made famous in his analysis of markets for lemons.

Goldfarb and Tucker (2019) systematize the digital difference as a reduction in five costs. Digital technology lowers *search* costs, making it easier for buyers and sellers to find one another and enabling markets - long-tail retail, global freelancing - that high search costs previously foreclosed. It lowers *replication* costs, as just discussed, to approximately zero. It lowers *transportation* costs for anything that can be rendered as bits, collapsing the economic

significance of distance for information goods and services. It lowers *tracking* costs, making it feasible to observe individual behaviour at scale and thus enabling personalization, targeted advertising, and the data economy examined in Chapter 6. And it lowers *verification* costs, facilitating the establishment of identity and reputation that underpins trust among strangers in online markets. Nearly every phenomenon treated in this book - from platform business models to algorithmic management to the crisis of journalism - can be traced to one or more of these cost reductions and to the strategic responses of firms, workers, and states.

1.4. Network Effects, Increasing Returns, and Platform Logic

If the cost structure of information explains why digital goods are priced strangely, network effects explain why digital markets concentrate. A good exhibits network effects - or, in the older vocabulary, network externalities - when its value to each user increases with the number of other users. The classic formulation is owed to Katz and Shapiro (1985), who distinguished *direct* effects, in which users value the network itself (the telephone, the messaging application, the social network), from *indirect* effects, in which users value the complements that a large user base attracts (the operating system valued for its applications, the games console valued for its games). Network effects create a distinctive competitive dynamic: markets tip. Once a network attains a decisive lead, its very size becomes its principal attraction, new users rationally join the largest network, and competition *for* the market replaces competition *in* the market.

W. Brian Arthur (1989) generalized the point in his analysis of increasing returns and lock-in: where adoption is self-reinforcing, small historical accidents can select which among competing technologies prevails, and the market may lock in to a technology that is not, in any timeless sense, the best. The QWERTY keyboard narrated by David (1985) became the emblematic, if contested, illustration. For the digital economy, the implication is that market structure is not the smooth reflection of relative efficiency assumed by textbook competition but a path-dependent outcome in which timing, expectations, and strategic manoeuvre - subsidizing adoption, courting complement developers, ensuring compatibility or strategically withholding it - play decisive roles (Shapiro & Varian, 1999).

The theory of *two-sided markets*, developed by Rochet and Tirole (2003) and Armstrong (2006), completed the analytical toolkit. Many digital businesses do not sell to a single class of customers but intermediate between two or more groups - riders and drivers, advertisers and audiences, developers and users - whose demands are interdependent. The platform's pricing problem is then not to set a price but to set a *price structure*: how much to charge each side,

and typically to subsidize the side whose participation the other side values most. This explains a pattern that mystifies everyday intuition: why so many of the most valuable services of the digital economy - search, social networking, email, maps - are offered to users at a price of zero. The user side is subsidized; the advertiser side pays; and the currency the user in fact tenders is attention and data, a point at which the economist's analysis of two-sided pricing meets the media scholar's analysis of the audience commodity (Smythe, 1977) and the critic's analysis of surveillance capitalism (Zuboff, 2019). Chapter 3 develops these mechanisms in detail; here it suffices to register that non-rival goods, near-zero marginal costs, and network effects jointly generate the characteristic morphology of the digital economy - a landscape of a few very large platforms surrounded by ecosystems of dependent complementors - and thereby set the agenda for the competition-policy debates examined in Chapter 10.

1.5. Measuring the Digital Economy and the Productivity Paradox

How large is the digital economy, and what has it done for prosperity? Both questions are harder than they appear. On the first, the measurement problem follows directly from the definitional problem of Section 1.1: national statistical systems, built to count physical output and market transactions, struggle with an economy whose characteristic products are intangible, whose characteristic prices are zero, and whose characteristic firm may book its profits in a jurisdiction other than the one where its users live. The OECD (2020) roadmap toward a common measurement framework recommends a tiered approach consonant with Bukht and Heeks (2017), supplemented by satellite accounts for digital activity; yet even sophisticated frameworks capture consumer surplus poorly. When a household substitutes a free navigation application for a purchased road atlas, measured GDP may fall while welfare rises. Brynjolfsson and collaborators have proposed measuring the consumer surplus of free digital goods directly, through massive online choice experiments eliciting what users would demand in compensation for giving up particular services - with results suggesting that conventional statistics omit welfare gains of macroeconomic magnitude (Brynjolfsson, Collis, & Eggers, 2019).

On the second question, the digital economy inherited a puzzle before it had a name. In 1987 Robert Solow famously quipped that "you can see the computer age everywhere but in the productivity statistics" (Solow, 1987, p. 36), crystallizing what became known as the productivity paradox: decades of heavy investment in information technology coincided, through the 1970s and 1980s, with a marked slowdown in measured productivity growth. Brynjolfsson (1993), surveying the evidence, catalogued candidate explanations - mismeasurement of output, lags in learning and organizational adjustment, redistribution of

gains among firms rather than net creation, and managerial mismanagement of the new technology. Paul David (1990) supplied the most influential historical analogy: the electric dynamo, whose productivity effects took some four decades to materialize, because they required not merely the installation of motors but the wholesale redesign of factories around the new power source. On this view, general purpose technologies deliver their gains only after complementary investments in organization, skills, and infrastructure - a lesson that recurred when United States productivity finally surged in the late 1990s, and that structures today's debates about the economic effects of artificial intelligence, where a second productivity paradox is widely discussed (Brynjolfsson, Rock, & Syverson, 2019).

The measurement debate is not an accountant's quarrel; it is where the epistemology of the digital economy is decided. What a society cannot measure it struggles to tax, to regulate, and to distribute. The sociologist will recognize in the statistics themselves an institution - a socially constructed apparatus of commensuration whose categories (what counts as a worker, a firm, an industry) the digital economy systematically strains, as the classification battles over gig workers examined in Chapter 7 will make vivid.

1.6. Digital Entrepreneurship: A First Approximation

Entrepreneurship theory begins, by common consent, with Joseph Schumpeter (1934), for whom the entrepreneur is the agent who carries out *new combinations* - new goods, new methods of production, new markets, new sources of supply, new forms of organization - and thereby, in the formulation of his later work, drives the process of creative destruction through which capitalism incessantly revolutionizes its own structure (Schumpeter, 1942). Alternative traditions emphasize the entrepreneur as the alert discoverer of profit opportunities (Kirzner, 1973) or as the bearer of uninsurable uncertainty (Knight, 1921). What, then, is *digital* about digital entrepreneurship - is it merely ordinary entrepreneurship that happens to use computers?

The contemporary literature answers no. Nambisan (2017), in an agenda-setting article, argues that digital technologies transform the very nature of entrepreneurial processes and outcomes in ways that existing theory only partly accommodates. Digital artefacts are *editable, reprogrammable, and recombinable*; consequently, the boundaries of the entrepreneurial product need never close, and the venture idea remains fluid long after launch - a condition institutionalized in the practices of continuous deployment, perpetual beta, and iterative pivoting celebrated by the lean startup movement (Ries, 2011). Digital infrastructures - cloud computing, app stores, payment systems, social media - dramatically lower the capital requirements of founding and allow ventures to rent, rather than build, the apparatus of scale,

compressing the time from foundation to global reach. And digital platforms redistribute entrepreneurial agency: the venture is frequently a complementor within an ecosystem governed by a platform whose rules - of access, data, and revenue sharing - condition its prospects as thoroughly as any market. Autio et al. (2018) extend the analysis to entrepreneurial ecosystems, arguing that digital affordances have shifted the locus of advantage from sector-specific knowledge spillovers to generic capabilities of business-model innovation, thereby reshaping where and how clusters of venturing form.

For the purposes of this book, digital entrepreneurship is defined as *the identification, evaluation, and exploitation of opportunities in which digital technologies are constitutive of the venture's product, of its process of creation, or of both*. The definition spans the spectrum from the venture-backed software startup to the platform-dependent micro-entrepreneur - the seller on a marketplace, the creator on a video platform, the developer in an app store - whose numbers vastly exceed the former and whose economic condition, poised ambiguously between enterprise and labour, will occupy us in Chapters 4, 5, and 7. It also signals a theme to which the sociological lens is indispensable: entrepreneurship is not merely an economic function but a social role, unevenly accessible across lines of class, gender, geography, and network position, and embedded in cultural narratives - the garage, the founder, the unicorn - that are themselves products of the media system.

1.7. Three Disciplinary Lenses: Economics, Sociology, Media Studies

It remains to make explicit the division of labour among the three disciplines to which this book is addressed. *Economics* contributes the theory of mechanism: the cost structures, incentive problems, and equilibrium tendencies sketched in Sections 1.3 and 1.4, together with the apparatus of measurement and the normative framework of welfare analysis on which competition policy and regulation draw (Shapiro & Varian, 1999; Goldfarb & Tucker, 2019). Without the economics, the observer of the digital world mistakes strategy for conspiracy and structure for accident; the zero price, the aggressive subsidy, the walled garden all have precise economic logics.

Lens	Core question	Typical mechanism	Primary risk
Economics	How are incentives, costs and prices structured?	Marginal cost, network effects, multi-sided pricing	Market power, inefficient allocation
Sociology	How are markets embedded in relations and inequalities?	Trust, networks, stratification, classification	Exclusion, cumulative advantage, opaque ranking

Media studies	How are visibility and public communication mediated?	Gatekeeping, attention allocation, platformization	Distorted public sphere, dependency, information disorder

Table 1.1. Three analytical lenses used throughout the book.

Sociology contributes the theory of embeddedness: the demonstration, classically stated by Granovetter (1985), that economic action is embedded in concrete networks of social relations, and that markets presuppose trust, norms, and institutions they do not themselves produce. The sociology of the digital economy examines how those relational preconditions are engineered online - reputation systems, ratings, verification - and with what stratifying consequences; how the network society reorganizes production around flexible, globally distributed arrangements (Castells, 2010); how digital inequalities of access, skill, and usage reproduce and refract older inequalities (DiMaggio et al., 2001); and how categories such as worker, firm, and market are contested and remade in the digital context. Without the sociology, the observer mistakes the platform economy for a frictionless market of atomized agents and misses the institutional work - and the power - that holds it together.

Media and communication studies contribute the theory of mediation: the recognition, running from McLuhan (1964) through the political economy of communication to contemporary platform studies, that the infrastructures of communication are never neutral conduits but shape what can be said, seen, and sold. The media lens discloses the digital economy's dominant firms as heirs of the broadcasting industry's audience-selling business (Smythe, 1977), now perfected through datafication and algorithmic curation; it analyses platformization as the penetration of platform logics into the production and circulation of culture (van Dijck et al., 2018); and it insists that the stakes of the digital economy include not only prices and productivity but the conditions of public communication on which democratic life depends. Without the media lens, the observer treats the attention economy as an ordinary market and misses that its raw material is the public sphere itself.

These lenses are not rivals but complements, and the chapters that follow will braid them deliberately. The platform analysed as a two-sided market in Chapter 3 returns as an infrastructure of sociality in Chapter 8 and as a curator of public discourse in Chapter 9; the entrepreneur theorized economically in Chapter 4 returns as a cultural figure and a precarious worker in Chapters 5 and 7. The wager of the book is that each return enriches, rather than repeats, the analysis.

WORKED EXAMPLE

Scenario. A free navigation service replaces a paid paper map costing EUR 20 per year for 500,000 users. Advertising revenue is EUR 3 million.

Calculation. Measured market revenue is EUR 3 million, although users may obtain services that previously required EUR 10 million of expenditure. GDP records the monetized transaction, not the full consumer surplus.

Interpretation. The example illustrates why digital welfare gains can coexist with modest measured output growth and why the productivity paradox is partly a measurement problem rather than a simple absence of value creation.

THREE LENSES

Economics asks: Which costs fall, which complements become scarce, and where do rents emerge?

Sociology asks: Which groups gain capabilities, status, or bargaining power, and which are excluded?

Media studies asks: Which infrastructures and interfaces organize visibility, attention, and cultural circulation?

1.8. Chapter Summary and Discussion Questions

This chapter has laid the conceptual and historical foundations of the volume. The digital economy was defined, following the tiered framework of Bukht and Heeks (2017) and the cost-based perspective of Goldfarb and Tucker (2019), as the ensemble of activities in which digitized data and connectivity are essential inputs to value creation - a definition that treats digital technology as a general purpose technology rather than a sector. Five historical waves were distinguished, from the mainframe era to the age of data-driven artificial intelligence, each layered upon its predecessors. The distinctive economics of information - high fixed and near-zero marginal costs, non-rivalry, experience-good character - was shown to underlie the peculiar pricing, bundling, and property arrangements of digital markets, while network effects, increasing returns, and two-sided market structures were shown to underlie their tendency toward concentration. The measurement of the digital economy and the productivity paradox were introduced as standing puzzles with high political stakes. Digital entrepreneurship was given a first definition emphasizing the fluidity, scalability, and platform-embeddedness of digital venturing. Finally, the three disciplinary lenses of the book - economic mechanism, sociological embeddedness, and media-theoretical mediation - were set out and their complementarity defended.

Discussion questions. (1) Take a firm of your choice and locate it within Bukht and Heeks's (2017) three tiers; what changes in your analysis if you move it from one tier to

another? (2) Solow's paradox held that computers were everywhere but in the productivity statistics; construct the strongest mismeasurement explanation and the strongest organizational-lag explanation, and propose evidence that would discriminate between them. (3) Why are so many digital services priced at zero, and in what senses is "free" a misnomer? Draw on the theory of two-sided markets and on the concept of the audience commodity. (4) Nambisan (2017) claims that digital technologies change the nature, and not merely the tools, of entrepreneurship; evaluate this claim against Schumpeter's (1934) classical account. (5) Identify a current controversy about a digital platform and show what the economic, sociological, and media-studies lenses each reveal - and each obscure.

1.8a. Numerical and Data Exercises

Exercise 1.A (network value). A messaging app has n users. Under Metcalfe's law the value of the network is proportional to the number of possible pairwise connections, $n(n-1)/2$. (a) If the network grows from 1,000 to 2,000 users, by what factor does its value rise? (b) Why does this convexity help explain "winner-take-most" dynamics? (c) Sarnoff's law (value $\propto n$) and Reed's law (value $\propto 2^n$) bracket Metcalfe's from below and above: which best fits a broadcast channel, and which a platform whose users form many overlapping groups?

Exercise 1.B (the productivity paradox). Suppose measured multifactor productivity growth averaged 1.5% per year in 1995-2004 but only 0.5% in 2005-2019, even as ICT investment rose. (a) State Solow's paradox in one sentence. (b) Give two "measurement" and two "implementation-lag" explanations. (c) If free digital services generate consumer surplus not captured in GDP, does mismeasurement bias measured productivity up or down?

Indicative answers. 1.A (a) $[2000 \cdot 1999/2] \div [1000 \cdot 999/2] \approx 4.0\times$ - value roughly quadruples when users double. (b) Each new user adds value in proportion to the existing base, so large networks pull ahead ever faster; combined with subsidies and expectations this can tip a market. (c) Sarnoff fits one-to-many broadcast; Reed fits group-forming social platforms. 1.B (a) "You can see the computer age everywhere but in the productivity statistics." (b) Measurement: free and quality-improved goods under-counted, intangible capital not capitalized; implementation lag: complementary organizational change and skills take time, and general-purpose technologies diffuse slowly. (c) Down - uncounted surplus means output, and hence productivity, is understated.

1.9. Case Studies and Debates

The following materials are designed for seminar use. Each case applies the chapter's concepts to a concrete episode; each debate states a genuinely contested question on which the literature divides, together with the strongest considerations on each side. Instructors may assign them as structured discussions, position papers, or examination questions.

Case 1.1. Measuring what is free: the GDP-B experiment

Section 1.5 argued that conventional national accounts systematically miss the welfare contribution of zero-price digital services: when a purchased road atlas is replaced by a free navigation application, GDP may fall while welfare rises. The research programme of Brynjolfsson and colleagues converts the complaint into measurement (Brynjolfsson, Collis, &

Eggers, 2019; Brynjolfsson, Collis, Diewert, et al., 2019). In massive online choice experiments, representative samples of users are confronted with incentive-compatible trade-offs - how much money would you require to give up search engines, email, digital maps, or social media for a month or a year? - and the elicited valuations are aggregated into an experimental welfare metric the authors call GDP-B, for the benefits that conventional GDP omits. The magnitudes are striking: median valuations for search and email run to thousands of dollars per user per year, orders of magnitude above the advertising revenue those users generate, implying consumer surplus of macroeconomic size invisible in the statistics on which growth debates rest.

The case rewards critical handling as much as admiration. Stated valuations, even incentive-compatible ones, measure willingness to accept the loss of a service *given the current organization of life around it*: the thousands of dollars demanded to abandon email measure, in part, the switching and coordination costs of Chapter 3, not pure welfare; the network effects that make the service valuable also make its hypothetical absence catastrophic, inflating the counterfactual. And a welfare metric is not a production metric: GDP answers questions - about taxable capacity, productivity, and debt sustainability - that GDP-B does not. Seminar questions: (1) Does the navigation-application example show that GDP is wrong, or that it is being asked the wrong question? (2) If free services were valued at elicited willingness to accept and added to output, what would happen to measured productivity - and to the productivity paradox of Section 1.5? (3) Who gains politically from larger measured digital welfare, and who from smaller?

Case 1.2. The dot-com crash as installation crisis

Between 1995 and 2000, the NASDAQ composite rose roughly five-fold; between March 2000 and October 2002 it lost some three-quarters of its value, extinguishing paper wealth measured in trillions and a generation of Internet ventures with it. Read through daily commentary, the episode was a verdict: the digital economy had been a mania, and the sober economy of atoms had reasserted itself. Read through Perez (2002), whose framework Section 1.2 introduced, the episode was instead a recurrence: every technological revolution since the canals, on her account, passes through an *installation* period in which financial capital, chasing the new paradigm's returns, decouples from production capital in a bubble whose collapse is the *turning point* - after which, under recomposed institutions, a *deployment* period spreads the

technology's gains through the wider economy. The prediction embedded in the framework is the one the following two decades vindicated: the crash destroyed the speculative superstructure while leaving the installed base - fibre-optic capacity laid at bubble prices, a trained cohort of engineers, consumer habits of being online - as the subsidized foundation on which the platform wave of the 2000s was built, largely by firms (founded in the ruins or surviving them) that acquired infrastructure and talent at post-crash prices.

Seminar questions: (1) Distinguish, in the dot-com episode, the claims that proved false (which ones, precisely?) from the timing that proved wrong. What does the distinction imply for evaluating present-day claims about artificial intelligence? (2) Perez holds that bubbles perform a function - over-funding infrastructure that rational finance would under-provide. State the strongest objection to this functionalism. (3) Identify one current technology market and argue, using observable indicators rather than intuition, whether it is in installation, turning point, or deployment.

Debate 1.1. Is 'the digital economy' still a useful category?

Motion: the term 'digital economy' should be retired, because it no longer designates a distinct object. *For the motion:* Section 1.1's own logic - the three-tier framework of Bukht and Heeks (2017) expanding until the broadest tier shades into the economy as such, and the general-purpose-technology view of Bresnahan and Trajtenberg (1995) - implies that 'digital economy' is now as informative as 'electricity economy' would have been in 1960: every sector qualifies, so the category discriminates nothing, and its persistence serves mainly the rhetorical interests of those who benefit from treating ordinary market power as novel and therefore unregulable. *Against the motion:* the category retains analytical bite precisely because the mechanisms this book assembles - non-rivalry, near-zero marginal cost, network effects, data feedback - apply with full force only to a describable subset of activity, and collapsing the distinction licenses the opposite rhetorical abuse: regulating supermarkets and social networks with the same instruments because 'everything is digital now'. The debate is not verbal: statistical classification (Section 1.5), the scope of gatekeeper regulation (Chapter 10), and research funding all turn on where the boundary is drawn, and by whom. Participants should be required to state, for their side, one policy decision that would be made differently if their position on the category were adopted.

1.10. In Depth: The Political Arithmetic of Measuring Digitalization

This advanced section develops the measurement problem of Section 1.5 at the level of detail a reader engaging the primary literature - or the policy debates - will require. Its organizing claim is Coyle's (2014): gross domestic product is not a natural fact awaiting better instruments but a construction, invented for the problems of depression and war mobilization, whose conventions embed decisions about what counts as production at all - and the digital economy is best understood as the largest shock those conventions have absorbed since the framework's invention.

Boundaries: what the accounts were built not to see

Three boundary conventions of the national accounts bear directly on digitalization. The first is the *production boundary*: unpaid household activity is, by convention, outside GDP - the classic textbook joke being that the economy shrinks when a person marries their housekeeper. Digitalization runs the joke at scale in both directions. When households substitute free self-service for purchased services - navigation for atlases and travel agents, search for reference publishing, platform photography for studios - measured output falls while activity continues; conversely, when platforms convert household assets and time into transactions - the spare room becomes lodging revenue, the commute becomes a ride - measured output rises without any new activity existing. The second convention is the treatment of *intangibles*: expenditures on software, data, and research hover ambiguously between consumed inputs and capital investment, and successive revisions that capitalized software and research each mechanically raised measured GDP - a reminder that growth statistics respond to accounting reform as well as to production. Data itself remains largely uncapitalized: the asset at the centre of Chapter 6 appears in no balance sheet the statistical system consolidates. The third convention is *zero-price output*: the accounts record market transactions, so services exchanged for attention and data (Chapters 3 and 5) enter GDP only through the advertising revenue of their producers - a sliver, as Case 1.1's valuations suggest, of their consumption value.

Prices and quality: the deflator problem

Real growth is nominal growth corrected by price indices, so everything depends on measuring price change for goods whose quality transforms continuously. The workhorse correction is hedonic pricing - regressing prices on characteristics (processing speed, storage, resolution) so that a constant-quality price can be tracked as the goods themselves mutate - and

its application to computing famously records some of the fastest sustained price declines in economic history. But hedonics requires enumerable characteristics, and the digital economy's frontier products resist enumeration: what is the characteristics vector of a search engine, a social network, or a model that writes? Where quality change goes unmeasured, inflation is overstated and real growth understated - the *mismeasurement hypothesis* in the productivity-paradox debate of Section 1.5. The careful empirical verdict, associated with the work synthesized in Brynjolfsson, Rock, and Syverson (2019), is that mismeasurement is real but insufficient: the productivity slowdown appears too large, too broad across sectors, and too internationally synchronized to be an artefact of missing consumer surplus in a handful of free services - which is why the implementation-lag explanation carries the burden in Chapter 6. The methodological lesson survives the verdict: every strong claim about digital-era growth or stagnation is, at bottom, a claim about deflators, and should be audited as such.

Satellite accounts and digital trade

The statistical community's constructive response, coordinated by the OECD (2020), is the *digital economy satellite account*: a supplementary presentation that leaves the core accounts intact while re-tabulating transactions along three digital dimensions - the *nature* of the product (digital or physical), the *mode* of transaction (digitally ordered, digitally delivered), and the *actors* involved, with digital intermediation platforms isolated as a category so that the platform economy of Chapter 3 becomes statistically visible rather than dissolved into retail, transport, and accommodation. The same framework organizes the measurement of *digital trade*, where the difficulties compound: digitally delivered services cross borders without customs events; the born-global firms of Chapter 4 book revenues in jurisdictions chosen for tax rather than activity, so that bilateral statistics record the geography of accounting more than of production; and free cross-border services - the search and social platforms consumed worldwide but sold, as advertising, elsewhere - constitute a trade flow the balance of payments barely registers. These are not technical curiosities: Chapter 10's digital-tax controversies and the development asymmetries of Section 10.3 are litigated on precisely these numbers, and the student who has worked through this section will recognize, in each headline statistic of that debate, the boundary conventions that produced it. Discussion questions: (1) Trace one convention - the production boundary, intangible capitalization, or the deflator - through a headline claim about the digital economy's size, and restate the claim with the convention made explicit. (2) Should data be capitalized in the national accounts? Specify the valuation method

you would use and the incentive effects of adopting it. (3) Design the three tables of a digital satellite account for your own country, and identify the data sources each would require.

1.11. In Depth: A Working Chronology of the Digital Economy, 1945-2025

Section 1.2's five waves compressed eight decades into a paragraph each; this section decompresses them into a working chronology of four periods whose boundaries deliberately cut across the waves - not a list of gadgets, but a narrative of the economic and institutional turning points on which the book's later arguments rest. Dates are anchors, not fetishes; the reader should attend to the *sequence of preconditions*, which is the chronology's argument.

1945-1975: computation becomes an industry

The stored-program computer emerged from wartime calculation and remained, for two decades, a capital good of states and giant firms - the mainframe era whose economics Section 1.2 sketched. The period's durable legacies are institutional as much as technical: the software contract was invented when hardware and programs were unbundled - a regulatory and commercial decision of the late 1960s that created software as a *product* and thus the industry Chapter 2's economics describes. The semiconductor learning curve was also established, with the planar process and the integrated circuit converting computation into a manufactured commodity whose price trajectory (Section 1.10's hedonic decline) became the era's master trend. Finally, the network was seeded - packet switching theorized and then built as a defence-funded research network in 1969, embedding at birth the end-to-end, open-protocol architecture whose political economy Chapter 9's infrastructure section examines. Post-industrial sociology (Bell, 1973) read the mainframe's professional class as the future; it was half right - the future belonged to information work, but not to the centralized machine.

1975-1995: decentralization and the software economy

The microprocessor (1971) made computation small enough to sell to persons, and the decade after 1975 built the personal-computer industry through the pattern Chapter 3 would later formalize: platform competition among incompatible systems, tipping toward the architecture that licensed widely - the openness-control trade-off of Case 3.2 played out in silicon - and the emergence of the packaged-software firm as the era's characteristic enterprise, selling the near-zero-marginal-cost good of Chapter 2 at scale. Two quieter developments completed the wave's preconditions for the next. The free-software movement (from 1983) codified, in licences, the commons production of Section 2.4 - an institutional invention as consequential as any product. And the research network outgrew its patrons: protocols standardized (1983), the network privatized by stages, and in 1989-1991 the World Wide Web layered a universal document system onto it - given away unpatented, a pricing decision (in

Chapter 2's terms) that determined the web would be a commons on which commerce would build rather than a product commerce would sell.

1995-2008: commercialization, crash, and the platform turn

The browser's commercialization (1994-1995) opened the Internet wave proper: electronic commerce and search founded (mid-1990s), the dot-com financing mania inflated and burst (Case 1.2), and the survivors consolidated the layers - search advertising's auction (Section 5.7) invented in the crash's shadow, turning attention into the auctionable commodity of Chapter 9. The platform turn arrived mid-decade: social networking scaled (from 2004), user-generated video (2005), and then the twin infrastructural announcements of 2006-2008 that this book treats as the fourth wave's infrastructural foundation - commodity cloud computing (Section 5.3), which collapsed the capital cost of founding (Chapter 4), and the smartphone with its application store (2007-2008), which put a sensor-laden, always-connected terminal in the world's pockets and created, in the app store, the archetypal governed marketplace of Sections 3.4 and 5.2. The financial crisis of 2008, meanwhile, set the monetary conditions - a decade of near-zero rates - whose venture-finance consequences Section 4.7 traces.

2008-2025: data capitalism and the intelligence wave

The smartphone decade fused the mechanisms this book assembles: mobile connectivity fed the data loops of Chapter 6, targeted advertising matured into the machinery of Section 5.7, the on-demand economy attached platform intermediation to physical services and labour (Chapters 5 and 7), and the platform constellation attained the market position that triggered the regulatory turn - data protection (2016-2018), the competition reports (2019), the gatekeeper and services acts (2022), chronicled in Chapter 10. The intelligence wave gathered beneath: the deep-learning demonstration of 2012 redirected the field, the transformer architecture (2017) made scale the strategy of Section 6.7, and the public deployment of conversational generative models (2022) converted a research trajectory into a macroeconomic expectation - with the investment boom, the training-data litigation (Case 6.2), and the first comprehensive AI statute (2024, Section 10.7) following within thirty months, the whole Perez sequence of Case 1.2 apparently restarting with the reader watching. The chronology's closing lesson is the one Section 1.2 promised: each wave's essential inputs - the chip's cost curve, the network's openness, the crash's infrastructure, the phone's sensors, the platform's data - were deposited by the wave before, and the analyst who would time the next one should watch, accordingly,

what the present wave is depositing. Discussion questions: (1) Identify, for each wave, the *pricing decision* (a good given away, unbundled, or auctioned) that most shaped the next wave. (2) The chronology names two institutional inventions - the software licence and the free-software licence - as turning points; defend the choice against a hardware-centred telling. (3) What is the intelligence wave depositing that its successor will presuppose?

Chapter 2. The Economics of Information and Digital Goods

Learning objectives. After studying this chapter, the reader should be able to: (1) characterize information as an economic good along the dimensions of rivalry and excludability, and explain Arrow's classic paradoxes of information markets; (2) derive the pricing consequences of a cost structure combining high fixed and near-zero marginal costs, including versioning, bundling, and the three degrees of price discrimination; (3) analyse how digitalization transforms each of the five costs identified by Goldfarb and Tucker - search, replication, transportation, tracking, and verification - and connect each transformation to an observable market phenomenon; (4) reconstruct the economic rationale for intellectual property and the principal empirical and theoretical challenges to that rationale; and (5) explain the economics of open-source production and commons-based alternatives to proprietary information goods.

KEY CONCEPTS

non-rivalry; excludability; fixed and marginal costs; versioning; price discrimination; search costs; verification costs; intellectual property; open models

Chapter 1 introduced, in outline, the claim that information has an economics of its own. The present chapter develops that claim systematically. It proceeds from the intrinsic properties of information as a good (Section 2.1), through the pricing strategies those properties compel (Section 2.2) and the cost transformations digitalization effects across markets (Section 2.3), to the institutional arrangements - intellectual property and its open alternatives - by which societies attempt to reconcile the production of information with its distribution (Section 2.4). Throughout, the exposition draws on a body of theory that predates the commercial Internet by decades: one of the quiet lessons of digital economics is how much of it was anticipated by the economics of information developed between the 1950s and the 1980s, before there was anything digital to apply it to.

2.1. Information as an Economic Good: Non-Rivalry and Excludability

Textbook economics classifies goods along two dimensions. A good is *rival* if one person's consumption diminishes the quantity available to others, and *excludable* if it is feasible, at reasonable cost, to prevent those who have not paid from consuming it. Private goods - bread, machinery, haircuts - are rival and excludable, and for them the price mechanism performs its celebrated coordinating work. Pure public goods - national defence, the classic

case analysed by Samuelson (1954) - are neither rival nor excludable, and for them decentralized markets characteristically fail, since no individual consumer has an incentive to pay for what she cannot be prevented from enjoying. Information in digital form occupies the awkward third cell of this taxonomy: it is intrinsically non-rival - a dataset, a song, or an algorithm can be used by any number of people simultaneously without depletion - while its excludability is not a natural fact but an institutional and technological achievement, produced by copyright law, encryption, paywalls, and licensing, and perpetually contested by the ease of copying.

The deepest early analysis of these difficulties is owed to Kenneth Arrow (1962), whose essay on the allocation of resources for invention identified a set of paradoxes that continue to structure the field. First, there is the *disclosure paradox*: the buyer of information cannot assess its value without inspecting it, but once she has inspected it she has, in effect, acquired it and need no longer pay - a radical form of the experience-good problem introduced in Chapter 1. Second, there is the *appropriability problem*: because information leaks, imitators free-ride on the costly discoveries of innovators, so that the private return to producing information falls short of its social return and markets, left to themselves, will underinvest in invention. Third, and pulling in the opposite direction, there is the problem of *efficient use*: once information exists, the socially optimal price of access is its marginal cost of reproduction - zero - so that any positive price excludes users who value the good above its true cost and thereby destroys welfare. The tension between the second and third points - between dynamic incentives to create and static efficiency in use - is the fundamental trade-off of information policy, and every institution examined in Section 2.4, from patents to open source, is best understood as a particular resolution of it.

The macroeconomic stakes of non-rivalry were drawn out a generation later by Paul Romer (1990), whose theory of endogenous growth placed ideas at the centre of the growth process precisely because of their peculiar properties as goods. Rival inputs - labour, capital, land - are subject to diminishing returns; ideas, being non-rival, can be used simultaneously across the entire economy, so that investment in the production of ideas yields increasing returns at the aggregate level. On this view, the digital economy represents something like the institutionalization of Romer's insight: an economy in which an ever-larger share of value consists of designs, software, media, and data - non-rival goods whose reproduction is costless - and in which, consequently, the classical scarcity assumptions of economic analysis apply to an ever-smaller share of what is produced. The abundance is real, but so is the paradox it

generates: an economy of goods that are costless to copy must construct scarcity - legal, technical, or attentional - in order to remain an economy of prices at all.

COMMON MISCONCEPTION

"Zero marginal cost" does not mean "zero economic cost" and does not imply that the efficient market price must be zero in every institutional setting. Digital products can carry large fixed creation, coordination, compliance, infrastructure, and discovery costs; pricing strategies are attempts to recover those costs under heterogeneous willingness to pay.

2.2. Zero Marginal Cost, Versioning, and Price Discrimination

How does a rational firm price a good whose marginal cost is zero? Not, evidently, at marginal cost: a price of zero recovers none of the fixed cost of creation, and an industry that priced this way would produce nothing. Nor, in general, at a single positive price: any uniform price excludes low-valuation consumers who could have been served at no additional cost, leaving welfare - and revenue - on the table. The economics of information goods therefore leads, with a kind of logical inevitability, to *price discrimination*: the practice of charging different prices for what is essentially the same good, calibrated to willingness to pay. The classical taxonomy descends from Pigou (1920). First-degree discrimination charges each buyer her exact reservation price; long a textbook curiosity, it has become a live prospect in the era of personalized pricing, when tracking data allow sellers to estimate individual willingness to pay with disquieting precision. Second-degree discrimination offers a menu of versions and lets buyers sort themselves; third-degree discrimination charges different prices to observable groups - students, seniors, geographic markets (Varian, 1989).

Shapiro and Varian (1999) showed that second-degree discrimination - *versioning* - is the natural strategy of the information economy. Because the marginal cost of any version is zero, the producer can create a product line at almost no expense: the delayed stock quote and the real-time feed, the free tier and the professional subscription, the standard resolution and the premium stream, the advertising-supported service and the paid ad-free variant. The art of versioning is to design the quality ladder so that high-valuation customers are not tempted by the cheap version - which sometimes requires deliberately *degrading* a product that would cost nothing to provide in full, a practice as economically rational as it is intuitively perverse. The freemium model that pervades software and media markets is versioning's contemporary apotheosis: the free tier serves simultaneously as a sampling device against the experience-good problem, a recruitment channel for network effects, and the bottom rung of a self-selection ladder.

Bundling is the second signature strategy. Bakos and Brynjolfsson (1999) demonstrated a result with far-reaching consequences: aggregating large numbers of information goods into a single bundle - the software suite, the streaming catalogue, the academic journal package - reduces the dispersion of consumers' average valuations, allowing the seller to extract more surplus with a single price than could be obtained by pricing each good separately. The logic depends on marginal costs being negligible, which is why bundles of unprecedented scale - tens of millions of songs, thousands of films - are peculiar to the digital economy. The subscription streaming model that reorganized the music, video, and publishing industries in the 2010s is, analytically, a giant bundle; its consequences for cultural production, for the remuneration of creators, and for the economics of media industries are examined in Chapters 5 and 9. Here the essential point is methodological: pricing patterns that appear novel, even predatory, to the casual observer - free products, degraded products, colossal bundles, personalized prices - fall out of a small set of principles once the cost structure of information is taken seriously.

2.3. The Five Cost Reductions: Search, Replication, Transportation, Tracking, Verification

Chapter 1 introduced Goldfarb and Tucker's (2019) organizing claim that digital economics is the study of what changes when five costs fall. The present section examines each in turn, connecting the abstract cost reduction to the concrete market phenomena it produces - and, importantly, to the new frictions and strategic behaviours that arise in response, for a recurring finding of the empirical literature is that lower costs do not simply produce the frictionless markets of early Internet prophecy.

Search costs. When finding a seller, a product, or a price is nearly free, several classical predictions follow: markets should approach the law of one price, obscure products should find their audiences, and intermediaries who profited from ignorance should disappear. The evidence is instructively mixed. Early comparisons of online and conventional retailers found that Internet prices were indeed lower, but that substantial price *dispersion* persisted online - identical books selling at meaningfully different prices across retailers - contrary to the frictionless benchmark (Brynjolfsson & Smith, 2000). Part of the explanation is strategic: sellers respond to price transparency by *obfuscating* - multiplying versions, add-on fees, and shipping schedules precisely to make comparison difficult, a dynamic documented econometrically by Ellison and Ellison (2009). The unambiguous success of low search costs lies elsewhere, in the *long tail*: the vast expansion of product variety available and findable online. Brynjolfsson et al. (2003) estimated that the consumer surplus generated by access to

previously unavailable niche titles exceeded, by a wide margin, the surplus generated by lower prices on popular ones - a finding popularized, with some exaggeration of the tail's commercial weight, by Anderson (2006).

Replication and transportation costs. The collapse of copying and delivery costs for anything expressible as bits is the material basis of everything discussed in Sections 2.1 and 2.2: it is what makes information goods effectively non-rival in practice and not merely in principle, and what detaches the economics of content from the economics of its former physical carriers - discs, paper, film stock. Two consequences deserve notice. Geographically, zero transportation cost for bits enables the global trade in digitally delivered services - software, design, analytics, media - that has allowed service exporters from Bangalore to Warsaw to integrate into world markets without container ships; the offshoring and remote-work phenomena examined in Chapter 7 are its labour-market face. Industrially, near-zero replication cost means that the binding constraint on cultural production shifts from distribution to attention: when everything can be delivered to everyone, the scarce resource is the audience's time, a point developed at length in Chapter 9. Waldfogel (2017) documents a striking upside: as digitalization slashed the costs of producing and distributing music, film, and books, the volume and, by several measures, the quality of new creative work rose - even as the industries' revenue models convulsed.

Tracking costs. The same infrastructure that delivers digital goods also observes their consumption. When every click, query, pause, and purchase can be recorded at negligible cost, personalization becomes feasible at scale: targeted advertising, recommendation systems, dynamic and individualized pricing, credit scoring on behavioural data. The economic logic is continuous with Section 2.2 - tracking is what makes first-degree price discrimination thinkable - but the phenomenon exceeds pricing. Cheap tracking converts the consumer's data trail into a factor of production, raising the questions of ownership, consent, and power that Zuboff (2019) has pressed under the heading of surveillance capitalism and that Chapter 6 treats in depth. It also creates a novel asymmetry: the seller may come to know the buyer's preferences better than the buyer knows the seller's costs, inverting the informational advantages assumed in classical market theory.

Verification costs. Markets among strangers require trust, and trust requires information about identity and past conduct. Digital systems lower the cost of establishing both, principally through reputation mechanisms - ratings, reviews, seller scores - that convert dispersed private experiences into a public signal. The economic significance of these mechanisms is difficult to overstate: they are the institutional foundation of marketplaces, gig platforms, and peer-to-peer

exchange, permitting transactions - boarding a stranger's car, renting a stranger's apartment - that verification costs previously confined to acquaintances or licensed professionals. Yet reputation systems are themselves economic objects with characteristic pathologies: inflated ratings, fear of retaliation in bilateral reviewing, fraudulent reviews, and cold-start exclusion of newcomers. Their sociological dimension - reputation as a new axis of stratification and discipline - is examined in Chapter 8.

2.4. Intellectual Property and Open Alternatives

Section 2.1 identified the fundamental trade-off: society wants information both created (which argues for rewarding creators with exclusive rights) and used (which argues for free access at zero marginal cost). Intellectual property - patent and copyright above all - is the oldest institutional resolution: a temporary, legally constructed monopoly that raises the private return to creation at the acknowledged cost of static inefficiency during the term of protection. The canonical economic analysis of copyright by Landes and Posner (1989) formalizes the balance: optimal protection trades the incentive effect of exclusivity against the welfare cost of supra-competitive pricing and the cost imposed on future creators, who build with materials the law has fenced. The qualification in the final clause is essential and easily overlooked: because creation is cumulative - every work borrows - stronger rights do not monotonically increase creation; beyond some point they tax it.

Digitalization destabilized this settlement from both directions. On one side, costless copying made infringement trivial, and content industries responded with the familiar arsenal: litigation, technical protection measures, term extensions, and a rhetoric of piracy. The empirical record complicated the rhetoric: in a widely debated study, Oberholzer-Gee and Strumpf (2007) found that music file-sharing's effect on record sales was statistically indistinguishable from zero, suggesting that downloads substituted mainly for purchases that would never have occurred - though subsequent studies using other markets and methods found significant displacement, and the question remains empirically contested. What is not contested is the institutional outcome: the industries stabilized not by defeating copying but by out-competing it, through licensed streaming bundles that made legal access more convenient than infringement - a resolution predicted, in effect, by the economics of Section 2.2. On the other side, critics argued that the digital environment exposed intellectual property's costs more than its necessity. Boldrin and Levine (2008) mounted the most radical case, contending that innovation historically flourished without, and often despite, intellectual monopoly; Lessig (2004) argued that expansive copyright, colliding with a medium built on copying, criminalized

ordinary creativity and locked up the raw materials of culture, and founded Creative Commons to build a licensing infrastructure for voluntary sharing.

The most consequential demonstration that alternatives are workable, however, came not from argument but from production. Free and open-source software - licensed to guarantee the freedoms to use, inspect, modify, and redistribute - grew from a programmers' ethic into critical infrastructure: the operating systems, web servers, databases, and machine-learning frameworks on which the commercial digital economy itself runs. Lerner and Tirole (2002) supplied the canonical economic account of why rational agents contribute without pay: deferred career benefits and signalling of skill to employers, intrinsic satisfaction, own-use value, and reputation within the community. Benkler (2006) generalized the phenomenon as *commons-based peer production* - a third mode of organizing production, alongside firms and markets, that non-rival goods and cheap coordination make viable at scale, with Wikipedia as its emblem beyond software. The lesson for economic theory is not that property is obsolete - open source rests, in fact, on a clever use of copyright through licensing - but that the space of institutions capable of eliciting the production of information goods is far wider than the patent-and-copyright settlement assumed, and that firms, states, and communities are still exploring it. The business models built atop open foundations, and the tensions between commercial platforms and the commons they enclose, recur in Chapters 5 and 6.

WORKED EXAMPLE

Scenario. A digital course costs EUR 120,000 to develop and EUR 0.50 per additional user to serve. A basic tier is priced at EUR 20 and a professional tier at EUR 80.

Calculation. If 4,000 users buy Basic and 1,000 buy Professional, revenue is EUR 160,000. Variable delivery cost for 5,000 users is EUR 2,500, leaving EUR 37,500 before marketing, support, taxes, and further fixed costs.

Interpretation. Versioning is not merely a marketing tactic: it is a mechanism for recovering high fixed costs when marginal replication costs are low.

THREE LENSES

Economics asks: How do non-rivalry, information asymmetry, and price discrimination shape market outcomes?

Sociology asks: Who can credibly signal quality, and how are trust and legitimacy distributed?

Media studies asks: How do formats, rights regimes, and distribution channels shape circulation and access?

2.5. Chapter Summary and Discussion Questions

This chapter has systematized the economics of information and digital goods. Information was located in the taxonomy of goods as intrinsically non-rival and only institutionally excludable, and Arrow's (1962) paradoxes - disclosure, appropriability, and the tension between dynamic incentives and static efficiency - were presented as the field's foundational problems, with Romer's (1990) growth theory drawing out the macroeconomic significance of non-rivalry. The pricing consequences of near-zero marginal cost were derived: value-based pricing, the three degrees of price discrimination, versioning as self-selection, freemium, and the large-scale bundling whose logic Bakos and Brynjolfsson (1999) established. The five cost reductions of Goldfarb and Tucker (2019) were examined individually, with attention to the counter-movements each provokes - obfuscation against search, attention scarcity against abundance, surveillance against tracking's conveniences, and reputation pathologies against cheap verification. Finally, intellectual property was analysed as one resolution among several of the fundamental trade-off, alongside the streaming settlement, the commons, and open-source production.

Discussion questions. (1) Reconstruct Arrow's disclosure paradox and identify three institutional devices - legal, technical, or reputational - by which contemporary markets for data and software mitigate it. (2) A software firm deliberately disables features in its free tier that cost nothing to provide. Is this welfare-reducing? Analyse using the theory of versioning and self-selection. (3) Why did price dispersion survive the Internet? Contrast the frictionless-market prediction with the evidence of Brynjolfsson and Smith (2000) and the strategic account of Ellison and Ellison (2009). (4) Evaluate the claim that streaming "solved" piracy. What does the resolution imply about the relative importance of price, convenience, and enforcement in markets for information goods? (5) Lerner and Tirole (2002) explain open-source contribution through career concerns and signalling; Benkler (2006) emphasizes intrinsic and social motivations. Design an empirical test that could discriminate between these accounts. (6) Should the copyright term be shorter for digital-native works? Argue both sides using Landes and Posner's (1989) framework.

2.5a. Numerical and Data Exercises

Exercise 2.A (versioning). A software firm faces 100 "professionals" (willing to pay up to EUR 200 for the full version, EUR 120 for a limited version) and 400 "casual" users (up to EUR 50 full, EUR 40 limited). Marginal cost ≈ 0. (a) What single full-version price maximizes revenue, and how much is it? (b) Design an incentive-compatible two-version menu that does better. (c) What does this show about zero-marginal-cost goods?

Exercise 2.B (bundling). Two goods A and B. Consumer 1 values (A = EUR 8, B = EUR 2); consumer 2 values (A = EUR 2, B = EUR 8). MC = 0. (a) Best separate prices and total revenue. (b) Best bundle price and revenue. (c) Why does bundling raise revenue here?

Indicative answers. 2.A (a) At EUR 50 all 500 buy → EUR 25,000, which beats EUR 200 (100 buyers → EUR 20,000). (b) Limited at EUR 40 (casual) and full at EUR 120: professionals' surplus from full (200−120 = 80) matches their surplus from the limited version (120−40 = 80), so they stay on full; revenue = 100×120 + 400×40 = EUR 28,000 > EUR 25,000. (c) With near-zero MC, quality is set to sort buyers, not to reflect cost; segmenting extracts more surplus than uniform pricing. 2.B (a) Each good priced at EUR 8 sells to its high-value buyer → EUR 16 total. (b) Both value the bundle at EUR 10, so a EUR 10 bundle sells to both → EUR 20. (c) Bundling compresses the dispersion of willingness-to-pay, letting the firm capture more of it.

2.6. Case Studies and Debates

The materials below apply the economics of information to two industry transformations and one live policy controversy. They presuppose Sections 2.1 through 2.4 and are suitable for seminar discussion or written assignment.

Case 2.1. From piracy to streaming: the recorded-music settlement

No industry lived the economics of this chapter more publicly than recorded music. Its twentieth-century model sold non-rival information bound to rival carriers - vinyl, cassette, compact disc - at prices sustained by control of manufacturing and distribution. Digital compression and peer-to-peer file-sharing, arriving together at the turn of the century, dissolved the carrier and with it the excludability the model presupposed: the marginal cost of a copy fell to zero for the consumer as well as the producer, and industry revenues roughly halved over the following decade. The enforcement response - litigation against services and individual users, technical protection, lobbying for stronger rights - is best read as an attempt to re-engineer

excludability by law and code. The empirical controversy over how much sharing actually displaced sales (Oberholzer-Gee & Strumpf, 2007, and the contrary literature noted in Section 2.4) mattered less to the outcome than a strategic fact: enforcement raised the cost of infringement without lowering the cost of compliance, and a price of zero with moderate inconvenience continued to defeat a positive price with different inconvenience.

The settlement, when it came, came from the other side of the ledger. Licensed streaming re-priced *convenience*: for a monthly fee - or free with advertising, in the freemium structure of Section 2.2 - the subscriber received the catalogue as a bundle (Bakos & Brynjolfsson, 1999, at industrial scale), instantly, everywhere, with search and recommendation attached. Legal access became easier than infringement, and industry revenues returned to growth on a wholly transformed base - ownership replaced by access, the album unbundled into the track and rebundled into the playlist, and the retailer's shelf replaced by an aggregator's ranking (Chapter 5). Waldfogel's (2017) verdict completes the case: measured by the volume and quality of new music, the digital era proved golden even while revenue convulsed - creation costs fell faster than revenues did. Seminar questions: (1) Allocate the industry's rescue among price, convenience, and enforcement; what evidence would change your allocation? (2) Who gained and who lost in the shift from ownership to access - distinguish labels, established artists, new artists, and listeners, and connect the artist-payment controversies to the bundle economics of Section 2.2. (3) Which industries could not replicate this settlement, and what property of their goods explains it?

Case 2.2. Open source as critical infrastructure - and the maintainer problem

Section 2.4 presented commons-based peer production as the demonstration that non-proprietary institutions can produce information goods of the first importance. The demonstration has since acquired a stress test. The open-source stack - operating-system kernels, web servers, cryptographic libraries, package ecosystems - now sits beneath essentially all commercial software, including the platforms and clouds of Chapters 3 and 5; yet much of it is maintained not by the ideal-typical crowd but by very small numbers of individuals, often unpaid, whose labour the depending firms consume as a free input. Eghbal (2020) documents the resulting *maintainer problem*: as usage scales, the maintainer's burden shifts from writing code to servicing demands - issues, reviews, support - that arrive from strangers at industrial volume; the non-rivalry of the code (Section 2.1) coexists with the acute rivalry of the maintainer's attention (Chapter 9's scarcity, arriving early); and periodic security crises in

ubiquitous, thinly maintained components have made the fragility a matter of infrastructure policy rather than community folklore.

The case complicates both celebratory and sceptical readings of the commons. Against the celebration: peer production does not escape the funding problem of Arrow (1962); it relocates it - from pricing the good to sustaining the producer - and the signalling incentives of Lerner and Tirole (2002) reward visible creation better than invisible maintenance. Against the scepticism: the response has been institutional innovation rather than enclosure - corporate sponsorship of maintainers, foundations stewarding critical projects, dedicated funds and security audits - which is this chapter's deepest lesson re-taught: the space of institutions for provisioning information goods is wide, and it is still being explored. Seminar questions: (1) Is the maintainer problem a market failure, a commons failure, or neither? Define your terms. (2) Firms free-ride on open source and also fund it; model the firm's decision and identify what tips it toward funding. (3) Should governments treat critical open-source components as public infrastructure? What follows practically if they do?

Debate 2.1. Personalized pricing: efficiency's ugly twin?

Motion: personalized (first-degree) price discrimination in digital markets should be legally restricted. Section 2.2 established the theory: with cheap tracking, sellers approach the textbook limit of charging each buyer her reservation price. *For the motion:* personalized pricing transfers essentially the entire surplus from consumers to sellers; it exploits informational asymmetry - the seller knows the buyer's willingness to pay while the buyer cannot observe the price others face - corroding the price system's public character; it prices desperation (insurance, credit, essential travel) most aggressively where elasticity is lowest; and it converts the privacy erosion of Chapter 6 directly into consumer harm, giving data protection an economic as well as dignitary rationale. *Against the motion:* the same mechanism that extracts surplus also *extends service* - under uniform pricing, low-valuation and often low-income buyers are excluded whom discrimination profitably serves; total output rises toward the efficient level as the deadweight loss of Section 2.1 shrinks; the objectionable cases (pricing desperation) are better addressed by sector rules than by banning a pricing technology; and enforcement would in practice prohibit discounts, loyalty pricing, and targeted promotions that consumers demonstrably value. The debate should be conducted with the welfare apparatus explicit: participants must state whose surplus counts, whether distribution or only efficiency matters, and what disclosure regime - if any - would change their vote.

2.7. In Depth: Search, Matching, and the Economics of Recommendation

Section 2.3 treated falling search costs as one item in Goldfarb and Tucker's (2019) five; this advanced section develops the underlying theory, because search is where the digital economy's most consequential intermediaries live and where the naive prediction - cheap search yields perfect markets - fails most instructively.

The Diamond paradox and the value of a little friction

The theoretical benchmark is startling and old. Diamond (1971) showed that in a market where consumers must pay even a *small* cost to observe each additional seller's price, the unique equilibrium is not the competitive price but the monopoly price: every seller, knowing that a marginal price cut cannot be observed without a search the consumer will not undertake for small stakes, prices as a monopolist, and the tiny friction propagates into maximal market power. The result - the Diamond paradox - overturns the comfortable intuition that markets degrade gracefully as frictions grow: the relationship between search costs and competition is discontinuous, and *who bears* the search cost, and *who controls* its size, become first-order strategic variables. Read forward, the paradox illuminates the digital economy twice. It explains the extraordinary value of search intermediaries: an actor that collapses the cost of the *next* price observation to zero detonates the Diamond equilibrium and forces sellers toward competition - which is why comparison infrastructure was among the web's earliest commercial layers, and why its neutrality matters so much. And it explains the strategic *reconstruction* of friction documented by Ellison and Ellison (2009) in Section 2.3: obfuscation - fee unbundling, version proliferation, drip pricing - is the seller's rational campaign to rebuild, inside the transaction, the search costs the infrastructure destroyed between transactions. The tug-of-war between friction-destroying intermediaries and friction-rebuilding sellers, not the frictionless equilibrium, is the permanent condition of digital markets.

From search to matching: two-sided frictions

Product search is the simple case: one side searches, the other is searched. The characteristic digital markets of Chapters 3, 5, and 7 - labour platforms, accommodation, dating, freelancing - are *matching* markets, in which both sides search, both sides choose, and a transaction requires mutual acceptance. Matching theory contributes three insights the search framework misses. First, *congestion*: when application is cheap, each attractive counterparty is flooded - the employer with hundreds of applications, the listing with dozens of requests - and the screening cost that cheap search saved reappears, multiplied, on the receiving side; platform

design answers with the rationing devices readers will recognize - application limits, response-rate scores, priority paid placement - each a priced re-imposition of scarcity. Second, *preference correlation*: where participants agree on who is desirable, matching concentrates on a short queue for the same counterparties (the superstar dynamics of Chapter 5's creator economy in miniature), whereas horizontal, differentiated preferences spread matches widely - so the *shape* of preferences, not only their intensity, determines how unequal a matching market becomes. Third, *thickness and timing*: matching markets can fail by being too thin (too few counterparties to search among) or effectively too thick (too many to evaluate), and the platform's management of when and how many options are shown is itself an allocation mechanism - the algorithmic dispatch of Chapter 7's ride-hailing being the limiting case in which the platform abolishes search entirely and simply assigns the match.

Recommendation: search inverted, demand shaped

The recommender system completes the trajectory: instead of the consumer searching the catalogue, the catalogue searches the consumer, using the tracking data of Section 2.3 to predict - in the vocabulary of Chapter 6 - which unseen item this person would value. The welfare economics is genuinely double-edged, and the empirical literature keeps both edges sharp. On one side, recommendation is the practical solution to the long tail's own paradox: Brynjolfsson et al.'s (2003) surplus from variety exists only if the variety is findable, and no unaided consumer searches millions of titles - discovery infrastructure is what converts abundance from warehouse fact into consumption value. On the other side, Fleder and Hosanagar (2009) demonstrated formally and by simulation that common collaborative-filtering designs - which recommend what similar users already consumed - contain a popularity feedback: they can *reduce* aggregate diversity, concentrating consumption on established hits even while each individual's consumption diversifies, because the system steers different users toward the same well-documented items and starves the undocumented tail it cannot yet evaluate. The design lesson generalizes far beyond retail: a recommender is not a neutral lens on given preferences but a demand-shaping institution whose objective function - predicted purchase, engagement time, diversity, revenue per impression - is chosen, and the choice allocates visibility exactly as Chapter 9's aggregators allocate attention. Discussion questions: (1) State the Diamond paradox precisely, and identify the assumption a price-comparison site breaks. (2) Choose a matching platform you know and inventory its congestion-management devices; who pays for each, in money or position? (3) Reconcile 'recommendation enables the long tail' with 'recommendation concentrates consumption' - which design parameters move a

system between the regimes? (4) If a marketplace's recommender maximizes platform revenue per impression, when do its recommendations diverge from the consumer's interest - and is that divergence detectable from outside (Chapter 10's audit problem)?

2.8. In Depth: Disclosure, Certification, and the Market for Quality

Sections 2.1 and 2.3 established that information goods are experience goods and that verification costs structure trust; this section develops the general theory of quality information - disclosure, certification, and ratings as economic institutions - following the survey of Dranove and Jin (2010), because the digital economy is, among other things, the largest quality-disclosure apparatus ever constructed, and its designers keep rediscovering, at scale, results this literature had already proved.

Unravelling and its failures

The theoretical benchmark is elegant: where disclosure is costless and verifiable, quality information should *unravel* - the best seller discloses to distinguish itself, whereupon the best of the remainder must disclose or be pooled with the worst, and so on down the ladder until silence itself is a confession and everyone discloses. The benchmark's value, as Dranove and Jin (2010) organize the evidence, lies in cataloguing why unravelling fails in practice, for each failure names a digital institution. Disclosure is costly and imperfectly verifiable - hence third-party *certification*, the outsourcing of verification to auditors, testing bodies, and platform badges (Chapter 8's 'top rated' thresholds are certification regimes in miniature). Consumers are inattentive or innumerate - hence the design of *formats*: coarse grades versus fine scores, a choice that trades information for usability and invites strategic bunching just above thresholds. And sellers respond to what is measured - the reactivity of Section 8.3 in its market form: schools teach to the ranking, hospitals select against risky patients, restaurants optimize the photographed dish; every certification scheme is a contract with Goodhart's law, and its design problem is to make gaming expensive relative to genuine quality improvement.

The certifier's own incentives

The literature's second movement turns the analysis on the certifier, and the digital economy supplies its largest cases. Certification is credence squared - the consumer must trust the certifier's diligence and independence - and the certifier's business model shapes both: who pays (the rated, the reader, or advertisers), whether the certifier competes for the rated party's custom, and whether its methodology is auditable. The pathologies are documented across domains - ratings shopping where issuers pay, grade inflation where certifiers compete on leniency, capture where the certified fund the certifier's revenue - and they transfer directly to the platform institutions this book has assembled: the marketplace that certifies its sellers while taxing their sales (Section 3.7's conflicted regulator), the app store whose review process is

both quality gate and competitive instrument (Debate 5.1), the search engine whose quality score co-determines advertising prices it collects (Section 5.7). The synthesis sharpens Chapter 8's conclusion into a design rule: quality information is a governed good, and the first question about any rating, badge, or ranking is the question this literature institutionalized - what does the certifier maximize, and who audits the auditor? Discussion questions: (1) State the unravelling argument and identify which of its assumptions each of three digital institutions - verified reviews, platform badges, app review - exists to repair. (2) Choose a certification scheme you rely on and analyse its payer structure; predict its characteristic bias. (3) Design a certification regime for AI-model capability claims (Section 10.7) that survives ratings shopping; who pays, in your design, and why?

Chapter 3. Platforms, Network Effects, and Two-Sided Markets

Learning objectives. After studying this chapter, the reader should be able to: (1) define direct and indirect network effects, distinguish them from economies of scale, and explain the critical-mass problem they generate; (2) present the theory of two-sided and multi-sided markets, including the logic of skewed price structures and the chicken-and-egg problem of platform launch; (3) analyse winner-take-most dynamics, specifying the conditions under which markets tip toward a single platform and the countervailing forces - multi-homing, differentiation, congestion - that preserve plurality; (4) explain lock-in and switching costs and evaluate the debate over whether network markets produce inefficient outcomes; (5) describe platform envelopment and the competitive dynamics of ecosystems, including the conflicted position of complementors; and (6) connect the economic theory of platforms to the sociological and media-theoretical analyses of platformization developed later in this volume.

KEY CONCEPTS

direct and indirect network effects; multi-sided markets; cross-side elasticity; tipping; switching costs; multi-homing; lock-in; envelopment; interoperability

If Chapter 2 explained why digital goods are priced strangely, this chapter explains why digital markets are structured strangely: why a handful of intermediaries - marketplaces, app stores, social networks, search engines, operating systems - stand at the centre of the digital economy, matching parties who could not efficiently find each other unaided, and extracting, in exchange, a share of the value of nearly every transaction they mediate. The platform is the characteristic firm of digital capitalism, as the factory was of industrial capitalism (Srnicek, 2017); understanding its economics is the precondition for every debate about its power.

3.1. Direct and Indirect Network Effects

The foundational observation predates the digital era. Rohlfs (1974), analysing demand for telecommunications services, formalized what every telephone company knew: the value of subscribing depends on who else subscribes, so that demand is interdependent rather than individual, and a service may be worthless below a *critical mass* of users yet immensely valuable above it. Katz and Shapiro (1985), as noted in Chapter 1, generalized the insight and drew the crucial distinction. *Direct* network effects operate within a single group: each additional user of a communication network, a social platform, or a compatible standard raises the value of participation for every other user. *Indirect* network effects operate through

complements: additional users of an operating system attract additional application developers, whose applications in turn attract additional users - a positive feedback loop running through a second market rather than within the first.

Popular discussion often compresses these dynamics into "Metcalfe's law," the engineering heuristic that a network's value grows with the square of its membership, since the number of possible connections does. The heuristic captures the qualitative point - value grows superlinearly in users - while overstating it: not all connections are equally valuable, and the marginal member of a mature network adds less than the marginal member of a growing one. More analytically consequential is the distinction between network effects and *economies of scale*. Both produce advantages of size, but scale economies operate on the cost side - average cost falls as output rises - whereas network effects operate on the demand side: the product itself becomes more valuable as adoption spreads. Digital platforms typically enjoy both simultaneously, along with a third reinforcing mechanism peculiar to the data economy: *data feedback loops*, in which more users generate more behavioural data, which improves the service through learning, which attracts more users (Cusumano et al., 2019). The stacking of these three mechanisms - supply-side scale, demand-side network effects, and data-driven learning - is what gives digital market leaders their formidable momentum.

Network effects also transform the strategic problem of the entrant. A conventional product must merely be better than its rivals; a network product must overcome the installed base - the accumulated value that incumbency itself confers. Below critical mass, even a superior network is inferior in use-value, which is why launch strategies in network markets so often involve aggressive subsidy, free service, celebrity seeding, or piggy-backing on an existing network's social graph. The critical-mass problem likewise explains the characteristic mortality pattern of platform ventures: most die not because their technology fails but because their network never ignites. These launch dynamics receive systematic treatment in the entrepreneurship chapters; here they serve to introduce the deepest version of the problem, to which we now turn - for most platforms must ignite not one network but two at once.

COMMON MISCONCEPTION

Network effects do not automatically produce monopoly. Their competitive consequences depend on multi-homing, differentiation, interoperability, congestion, local versus global network structure, switching costs, and the ability of entrants to mobilize one side of the market.

3.2. Two-Sided and Multi-Sided Markets

A platform, in the economic sense, is an intermediary that creates value by enabling direct interaction between two or more distinct groups whose demands are linked by indirect network effects (Rochet & Tirole, 2003; Evans & Schmalensee, 2016). The marketplace links buyers and sellers; the ride-hailing service links riders and drivers; the operating system links users and developers; the advertising-funded medium links audiences and advertisers; the payment card links cardholders and merchants - the example from which Rochet and Tirole's theory originally grew. What distinguishes the platform from the ordinary reseller is that the platform does not take possession of what is exchanged; it sells *access and matching*, and its product is, in a precise sense, the other side of the market.

Multi-sided platform: value and payment move on different paths

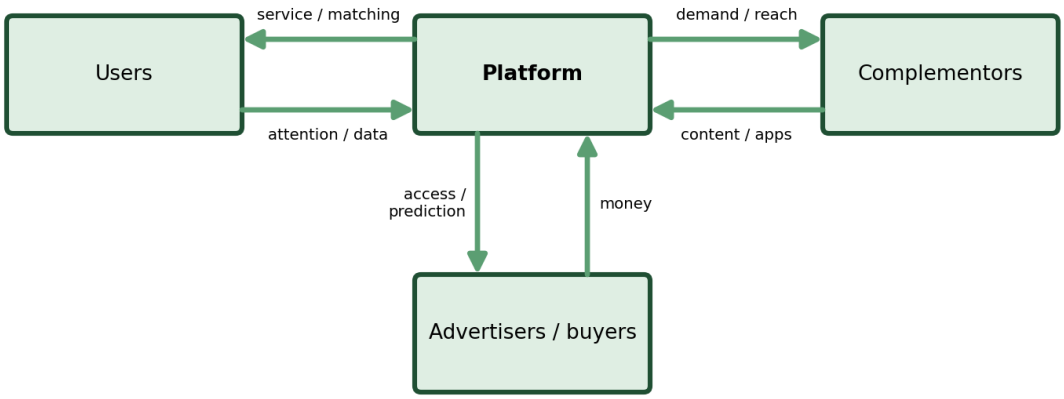

Figure 3.1. Multi-sided platform: cross-side value, subsidy and payment flows.

The theory's central result concerns pricing. Because the two sides' demands are interdependent, the profit-maximizing platform does not price each side by its own cost and elasticity, as a conventional multiproduct firm would; it chooses a *price structure* that internalizes the cross-side externalities (Rochet & Tirole, 2003; Armstrong, 2006; Parker & Van Alstyne, 2005). The side whose participation the other side values most - or whose own demand is most price-sensitive - is charged little, nothing, or less than nothing (free services, sign-up bonuses, subsidized hardware), while the other side pays above cost. Hence the durable asymmetries of the digital economy: audiences ride free while advertisers pay; cardholders earn rewards while merchants pay interchange fees; developers pay commissions while basic users pay nothing. Eisenmann et al. (2006) codified the pattern for strategists as the choice of a *subsidy side* and a *money side* - a vocabulary that has since passed into managerial common

sense. The normative implication, developed in Chapter 10, is that competition analysis in two-sided markets cannot proceed side by side: a price below cost on one side is not, by itself, predation, and a price above cost on the other is not, by itself, exploitation, since only the structure as a whole is economically meaningful.

The second canonical problem is *launch*. Caillaud and Jullien (2003) gave the chicken-and-egg problem its formal statement: neither side will join a platform without the other already present. The strategic repertoire that has evolved in response is by now well catalogued (Parker et al., 2016; Evans & Schmalensee, 2016): subsidize the more price-sensitive side, sometimes permanently, and recruit marquee participants whose presence disproportionately attracts the other side. Alternatively, solve one side yourself, as when a marketplace seeds its catalogue with its own inventory or a content platform commissions its own programming. Or begin as a single-sided product that accumulates one group before opening to the second, or constrain launch to a narrow niche - a single campus, city, or category - where critical mass is attainable, and expand concentrically from the beachhead. The launch problem also illuminates a structural feature of platform finance: because the returns to ignition are enormous and the probability small, platform ventures burn extraordinary quantities of capital in subsidies during the race for critical mass - a logic of venture-financed land-grabbing whose consequences for entrepreneurship and for labour markets are examined in Chapters 4 and 7.

Multi-sidedness, finally, is a matter of degree and design rather than a fixed attribute. Platforms add sides as they mature - the social network adds advertisers, then developers, then merchants; the marketplace adds logistics clients and advertising customers - and each added side multiplies both the cross-externalities the platform can internalize and the conflicts of interest it must manage. The platform that sells advertising against its users' attention while selling those same users subscriptions, and that ranks in its search results the merchants with whom it also competes, is not an aberration but the mature form of the species; the governance problems this form generates occupy Sections 3.3 and 3.4.

3.3. Winner-Take-Most Dynamics, Lock-In, and Switching Costs

Do network markets inevitably collapse into monopoly? The theoretical mechanism of tipping, introduced in Chapter 1, suggests they might: positive feedback rewards the leader, expectations converge on the likely winner, and the market selects a single platform. Yet the empirical landscape is more varied - some markets are dominated by one firm, others sustain durable oligopolies, and still others remain fragmented. The analytical task is therefore to specify the *conditions* of tipping. Eisenmann et al. (2006) identify three: markets tip toward a

single platform when network effects are strong, when users' needs are homogeneous (so that one platform can serve everyone), and when *multi-homing* - simultaneous participation on several platforms - is costly. Conversely, cheap multi-homing and heterogeneous demand preserve plurality: merchants list on several marketplaces, advertisers buy across several networks, drivers run two ride-hailing applications at once, and differentiated platforms coexist by serving distinct niches. Congestion and negative network effects - crowding, spam, declining match quality at scale - impose further limits. The precise phrase, then, is *winner-take-most* rather than winner-take-all: concentration is the tendency, monopoly the limiting case.

Where tipping occurs, its durability rests on *switching costs* and *lock-in*. Klemperer (1987) established the general theory: when changing suppliers imposes costs on the user - learning a new system, re-entering data, forfeiting accumulated benefits, coordinating the move with others - competition for the customer front-loads (firms fight fiercely for first-time users, then harvest the installed base), and market shares ossify. Farrell and Klemperer (2007), synthesizing two decades of research, showed how switching costs and network effects compound: the individual user of a dominant platform is deterred not only by her private switching cost but by the coordination problem of moving her whole network - contacts, followers, reviews, reputational capital - which no individual can move alone. Digital platforms engineer both layers deliberately: proprietary data formats, non-portable reputations and purchase histories, bundled ecosystems of devices and services, and loyalty schemes are, analytically, investments in the height of the exit barrier. The policy response - mandated data portability and interoperability - aims squarely at these mechanisms and is examined in Chapter 10.

Whether lock-in produces *inefficiency* is a separate and more contested question. The strong claim, descending from Arthur (1989) and David (1985), is that path dependence can trap markets on inferior standards - QWERTY over rival keyboard layouts, VHS over Betamax - because no individual adopter can profitably deviate from the installed base. Liebowitz and Margolis (1994) mounted the influential rebuttal: on their re-examination, the canonical examples of harmful lock-in largely dissolve - the alleged superiority of the alternatives was weaker than legend held, and where a genuinely superior technology offered gains sufficient to repay the costs of transition, markets did in fact switch. The debate matters because it fixes the burden of proof for intervention: if markets escape inferior standards on their own, the case for regulatory correction of tipping is weak; if coordination failures are real and persistent, waiting for self-correction cedes the market to the incumbent. The honest summary is that theory establishes the possibility of inefficient lock-in, that clean historical demonstrations are rarer

than the theory's popularity suggests, and that the modern platform case differs from the typewriter case in one crucial respect: today's incumbents actively manage their lock-in with a sophistication no keyboard standard ever possessed.

3.4. Platform Envelopment and Ecosystem Competition

Competition among platforms rarely takes the textbook form of similar firms contesting a well-defined market. Its characteristic manoeuvre is instead *envelopment*: a platform enters an adjacent platform's market by bundling the target's functionality with its own, leveraging a shared user relationship and the data that relationship yields (Eisenmann et al., 2011). The operating system envelops the browser, the media player, the office suite; the social network envelops messaging, photo-sharing, news distribution, marketplace classifieds; the marketplace envelops payments, logistics, advertising, cloud services. Envelopment succeeds where standalone entry would fail because the attacker does not need to solve the chicken-and-egg problem afresh - it already possesses one side of the market - and because bundling converts the incumbent's product into a feature. The strategic consequence is that platform markets have unstable boundaries: the relevant competitive arena is not the product category but the overlapping user base, and today's complement is tomorrow's rival.

This instability defines the predicament of the *complementor* - the developer, merchant, or creator who builds on a platform she does not control. The ecosystem literature (Gawer & Cusumano, 2014; Parker et al., 2016) frames the relationship as symbiotic: the platform supplies infrastructure, distribution, and trust; complementors supply the variety and innovation that no single firm could generate, and openness to them is a strategic variable the platform tunes - too closed, and the ecosystem starves; too open, and quality and coherence erode (Boudreau, 2010). But symbiosis coexists with structural asymmetry. The platform sets and revises the rules - commissions, ranking algorithms, data access, permissible categories - unilaterally; it observes every complementor's sales data; and it can integrate into any complement whose success that data reveals. Zhu and Liu (2018), studying a large marketplace econometrically, found that the platform's own entry systematically targeted the product spaces of its most successful third-party sellers - evidence that the fear of expropriation under which complementors operate is not paranoia but a rational reading of incentives. The result is a distinctive form of dependence that scholars of work and media will recognize from other chapters of this book: the app developer, the marketplace seller, and the platform creator are formally independent entrepreneurs whose fortunes are governed by private rule-makers - a condition to which the sociological concept of platform dependence and the legal debates of Chapter 10 both respond.

It is at this point that the economic analysis hands over, without discontinuity, to the sociological and media-theoretical analyses developed later in the volume. What economics describes as ecosystem governance, van Dijck et al. (2018) describe as the platform society: the extension of platform mechanisms - datafication, commodification, algorithmic selection - into sector after sector of social life, from news to urban transport to education and health. The two descriptions are not rivals; they are the same object viewed at different magnifications. The economist explains why the intermediary position is so profitable and so defensible; the sociologist and the media scholar explain what it means for a society when its markets, its public communication, and increasingly its infrastructures are governed from that position.

WORKED EXAMPLE

Scenario. A marketplace charges sellers 10% and buyers zero. At 20,000 monthly transactions with an average value of EUR 40, gross merchandise value is EUR 800,000.

Calculation. Platform revenue is EUR 80,000. If buyer acquisition costs EUR 4 per active buyer and 8,000 buyers are needed, acquisition expenditure is EUR 32,000 before trust, payments, support, and infrastructure costs.

Interpretation. The zero price on one side is not evidence of a non-commercial service; it may be the profit-maximizing price structure when buyer participation raises seller willingness to pay.

THREE LENSES

Economics asks: Which cross-side externalities determine the optimal price structure?

Sociology asks: How do trust, norms, and reputation stabilize exchange among strangers?

Media studies asks: How does platform governance determine discoverability, ranking, and access to audiences?

3.5. Chapter Summary and Discussion Questions

This chapter has assembled the theory of platforms. Network effects were traced from Rohlfs's (1974) interdependent demand through Katz and Shapiro's (1985) distinction between direct and indirect effects, and were differentiated from economies of scale and augmented by data feedback loops. The theory of two-sided markets (Rochet & Tirole, 2003; Armstrong, 2006) was presented with its two canonical results: the skewed price structure that internalizes cross-side externalities, and the chicken-and-egg launch problem (Caillaud & Jullien, 2003) with its repertoire of ignition strategies. Winner-take-most dynamics were analysed via the tipping conditions of Eisenmann et al. (2006), the switching-cost theory of Klemperer (1987) and Farrell and Klemperer (2007), and the efficiency debate between the path-dependence tradition and its critics (Liebowitz & Margolis, 1994). Finally, envelopment (Eisenmann et al.,

2011) and ecosystem competition were examined, with the complementor's structural dependence - documented empirically by Zhu and Liu (2018) - serving as the bridge from platform economics to the platform society.

Discussion questions. (1) Take a platform you use daily and map its sides, its subsidy structure, and the cross-side externalities each side exerts on the others. Which side is the money side, and why that one? (2) Ride-hailing drivers commonly multi-home; social network users rarely do. Using the tipping conditions of Section 3.3, explain the difference and predict which market structure each implies. (3) "A zero price on one side of a two-sided market tells us nothing about market power." Defend and attack this proposition. (4) Reconstruct the Liebowitz-Margolis critique of the QWERTY story. Does their argument transfer to contemporary platform lock-in, or do data, reputation, and cloud ecosystems change the analysis? (5) You advise a successful marketplace seller whose category the platform has just entered, replicating the seller's best-selling product. What strategic options remain, and what does each assume about the platform's incentives? (6) In what precise sense is 'platform capitalism' (Srnicek, 2017) an economic claim, and in what sense a sociological one? Use the envelopment and ecosystem material of Section 3.4 to make the distinction concrete.

3.5a. Numerical and Data Exercises

Exercise 3.A (two-sided pricing). A platform intermediates two groups whose demands are interdependent: each side values the other's participation. (a) Explain why the profit-maximizing price on one side can be below cost, or even negative. (b) If the platform can charge only one side, which should it subsidize, and on what does the answer depend? (c) Define "tipping" and give one condition that makes it more likely.

Exercise 3.B (switching costs). A customer pays EUR 30/month to an incumbent SaaS vendor. A rival offers the same service at EUR 22/month, but switching costs a one-off EUR 200. (a) Over what horizon does switching pay off, ignoring discounting? (b) If the customer discounts future savings at 1%/month, does the break-even lengthen or shorten? (c) How can the incumbent exploit this?

Indicative answers. 3.A (a) Because a subsidized side attracts the other, more valuable side; the loss on one side is recovered on the other. (b) Subsidize the side that the other values more, and/or the more price-sensitive side; the answer depends on the relative strength of the cross-side externalities and on each side's own elasticity. (c) Tipping = the market collapsing toward one platform; more likely with strong positive feedback (large network effects) and high multi-homing/switching costs. 3.B (a) Saving = EUR 8/month; EUR 200 ÷ EUR 8 = 25 months. (b) It lengthens - discounting lowers the present value of future savings, so more than 25 months are needed. (c) Raise switching costs (proprietary data formats, long contracts) or offer a targeted loyalty discount just large enough to erase the rival's net advantage.

3.6. Case Studies and Debates

The materials below exercise the platform toolkit of this chapter on three formative episodes and one standing policy controversy. They are designed to be worked with the concepts named explicitly: tipping conditions, multi-homing, switching costs, envelopment, and the complementor's predicament.

Case 3.1. The browser wars: envelopment before the concept

The first great antitrust confrontation of the Internet era supplied, avant la lettre, the textbook demonstration of Section 3.4. In the mid-1990s the browser - the universal client through which users reached the web - threatened to become a platform in its own right: paired with cross-platform runtime technologies, it promised to make applications independent of the underlying operating system, commoditizing the layer on which the incumbent's power rested.

The incumbent's response was envelopment executed with the full repertoire this chapter has assembled: the challenger's functionality was bundled with the dominant operating system at a price of zero, distribution agreements with computer manufacturers and access providers foreclosed the rival's channels, and the challenger's browser share collapsed from dominance to marginality within a few years. The ensuing litigation - *United States v. Microsoft Corp.* (2001) - produced findings that read as annotations to Sections 3.1 through 3.4: the court accepted that network effects and the 'applications barrier to entry' (an installed base of complements deterring operating-system entry - indirect network effects under a forensic name) sustained durable monopoly, and that bundling and exclusive dealing had unlawfully maintained it, while the remedy - conduct restrictions rather than the initially ordered breakup - set both the template for platform antitrust since and the pattern of disappointment with its remedies.

Seminar questions: (1) Reconstruct the episode in this chapter's vocabulary: identify the two-sided structures, the tipping conditions, and the envelopment move. Which element was decisive? (2) The zero price of the bundled browser was central to the case. Using Section 3.2, explain why 'below cost' is analytically treacherous in platform markets - and why the court found harm anyway. (3) The bundled browser itself later lost leadership to an entrant sponsored by a search platform. Does this vindicate Liebowitz and Margolis's (1994) scepticism about durable lock-in, or does it illustrate that only platforms dethrone platforms? What evidence separates the readings?

Case 3.2. Standards wars: VHS, Betamax, and the limits of the lock-in legend

The videocassette format war of the late 1970s and 1980s is the second canonical episode through which economists argue about Section 3.3. The stylized legend holds that the technically superior format lost: Betamax, first to market and backed by a formidable innovator, was defeated by VHS, and the outcome is cited alongside QWERTY (David, 1985) as proof that network markets select badly. The revisionist reading - in the spirit of Liebowitz and Margolis (1994) - retells the same facts as competition working: the formats' picture quality was comparable, but VHS offered the longer recording time that consumers demonstrably wanted (a full feature film, an American football broadcast). Its sponsor licensed the standard widely while the rival kept control tightly, so VHS machines were cheaper and more available. The pivotal indirect network effect - video-rental libraries stocking the more common format - amplified a lead earned on preference, not on accident. On this reading the market tipped, as Section 3.3 predicts under strong network effects, homogeneous needs, and costly multi-

homing (few households bought both machines) - but it tipped toward the product more consumers preferred, and the ‘inferior winner’ is legend, not finding.

The case earns its place in a platform chapter because the strategic anatomy - openness against control - recurs in every subsequent standards contest, from operating systems to charging connectors to AI model ecosystems. The sponsor who licenses widely trades margin for adoption, mobilizing complementors (Section 3.4) as allies in the tipping race; the sponsor who keeps control banks on quality and captures more value per unit if victory comes. Boudreau’s (2010) evidence on the innovation effects of granting access versus devolving control gives the trade-off its modern empirical form. Seminar questions: (1) Specify, for the videocassette war, each of Eisenmann et al.’s (2006) three tipping conditions, with evidence. (2) Construct the strongest surviving version of the ‘inefficient lock-in’ claim after the revisionist critique - what would have to be true, and how would a researcher detect it? (3) Choose a current standards contest and predict its outcome using the openness-control trade-off; state the observation that would falsify your prediction.

Case 3.3. BLIK: igniting a payment network from Central Europe

Section 3.2 argued that the hardest problem of a two-sided platform is ignition. Poland’s mobile payment system BLIK shows how that problem can be solved by a coalition of incumbents rather than by a venture-backed entrant. Launched in 2015 by Polski Standard Płatności, a company set up by several of the country’s largest banks and later joined by Mastercard as a shareholder, BLIK lets a user authorize an online or in-store payment, a cash withdrawal or a transfer to a phone number by entering a six-digit code generated in her own banking application. By the end of 2025 it counted 20.7 million active accounts in a country of roughly 37 million people, and in that year alone it processed some 2.9 billion transactions worth PLN 441.5 billion, almost half of them in e-commerce (Polski Standard Płatności, 2026).

Read with this chapter’s toolkit, the case is instructive on three counts. First, ignition: the consortium solved the user side by itself in the sense of Section 3.2, because BLIK was embedded in banking applications that millions of customers already used, so no separate sign-up was required. Merchants and payment intermediaries therefore faced a user base that existed on the first day. Second, the price structure: the user side pays nothing, while the merchant side pays for acceptance - the skew of Section 3.2. The competitive benchmark, however, was card payments and bank-transfer ‘pay-by-link’ services, against which BLIK competed on convenience and security as much as on price. Third, multi-homing and tipping: users continue to hold cards, and most merchants accept both, so the market did not tip to a monopoly. Yet in

Polish e-commerce BLIK became a default checkout choice, which shows how a locally rooted network can hold its own against global card schemes when it controls an interface that customers already trust.

The case also has a sovereignty dimension that anticipates Section 10.3. European policymakers have long regretted the dependence of card payments on non-European schemes. BLIK now operates in Slovakia and Romania as well as Poland, and its interoperability pilots with Spain's Bizum and Portugal's MB WAY within the EuroPA association make it a candidate building block of a European payments layer (Polski Standard Płatności, 2026). Seminar questions: (1) Map BLIK's sides and identify which side was subsidized during ignition and which is the money side today. (2) Why did competing banks cooperate on a shared network, and what governance problems inside such a consortium should the ecosystem literature of Section 3.4 lead you to expect? (3) Compare BLIK with M-Pesa (Case 4.1) and India's UPI (Section 4.5a): which ignition strategy did each use, and which actor - operator, bank consortium or state - carried the launch cost?

Debate 3.1. Remedying platform power: breakup, interoperability, or forbearance?

Motion: mandated interoperability, not structural breakup, is the correct remedy for entrenched platform power. *For the motion:* the sources of platform dominance identified in this chapter - network effects, switching costs, data feedback - are demand-side and would survive dismemberment: breaking a social network into pieces destroys user value while the pieces race to re-tip, whereas interoperability attacks the mechanism itself, letting users reach counterparties across platforms and carry identities, graphs, and reputations with them, so that competition occurs *in* the market rather than only *for* it; telecommunications interconnection is the working precedent. *Against the motion:* interoperability is regulation-intensive in exactly the way ex ante regimes (Chapter 10) are criticized for being - standards must be specified, updated, and policed against strategic degradation by the very incumbent they discipline, creating a permanent supervisory apparatus and freezing product architectures. Security and privacy furnish both real costs and inexhaustible pretexts. And where dominance was acquired or is maintained by acquisition and self-preferencing, structural separation is self-executing in a way conduct rules never are - the firm cannot disobey a divestiture. *Forbearance* - the position that platform leadership is serially contestable and each wave (Chapter 1) unseats the last - completes the triangle and should be assigned its own advocate. Participants must tie each remedy to the specific mechanism it addresses, name one historical deployment of it, and identify the failure mode their own position most fears.

3.7. In Depth: The Platform as Regulator - Governance as Institutional Design

Sections 3.2 through 3.4 analysed the platform's pricing and competitive strategy; this advanced section analyses its *legislation*. The framing is owed to Boudreau and Hagiu (2009): multi-sided platforms are best understood as private regulators, deploying instruments far richer than price - access rules, quality standards, information disclosure, dispute procedures, technical constraints - to solve the market failures that would otherwise destroy their ecosystems. The section develops the regulator's toolkit, the failures it answers, and the accountability problem it creates; it is the analytical bridge between this chapter's economics and the public regulation of Chapter 10.

Why price is not enough: the platform's market failures

A platform hosting thousands of independent complementors and millions of users internalizes, in miniature, the classic catalogue of market failures. *Externalities among participants*: one seller's counterfeit degrades buyers' trust in every seller (the lemons dynamic of Sections 1.3 and 2.8 operating on the platform's collective reputation); one developer's malware taxes the whole ecosystem; one abusive user poisons a forum. Because each participant bears only a sliver of the damage it causes to the common asset - the platform's trustworthiness - price alone cannot allocate correctly, and the platform regulates: prohibitions, review gates, certification tiers. *Information asymmetries*: buyers cannot assess sellers, so the platform mandates disclosure and operates the reputation machinery of Chapters 2 and 8 - legislated transparency, not emergent custom. *Coordination and standards*: complementors need stable interfaces and compatible expectations, so the platform issues technical law - APIs, design guidelines, deprecation schedules - whose revisions redistribute fortunes across the ecosystem as surely as any statute. The regulatory reading dissolves a false dichotomy that popular commentary keeps alive: the choice was never between the governed platform and the free one, because an ungoverned multi-sided market unravels; the choice is among governance designs, and among governors.

The toolkit: access, information, and dispute resolution

The platform-regulator's instruments repay classification. *Access control* runs from open registration through vetting to exclusive licensing, and Boudreau's (2010) evidence gives the trade-off its empirical form: granting broader access multiplies innovation by multiplying experimenters, while devolving more control to complementors deepens each one's investment - but openness beyond a threshold crowds the ecosystem and erodes average quality, so optimal

openness is interior and varies by layer. *Information design* determines what each side sees: which metrics are published (response rates, on-time scores), which are withheld (the exact ranking formula, per Section 7.1's information asymmetries), and which are engineered (badges, 'top rated' thresholds that function as regulatory categories, changing behaviour the way any certification does). *Sanctioning and dispute resolution* completes the state analogy: platforms operate notice, suspension, and termination regimes; adjudicate buyer-seller disputes with remedies (refunds, account restoration) and appeals of varying depth; and do so at a caseload no public court system approaches - a mass administrative justice, privately run, whose due-process quality Chapters 7 and 9 find wanting and whose reform is precisely what the Digital Services Act's procedural obligations (Chapter 10) attempt to compel. The comparative-institutional question is therefore not whether platform governance exists but how its design compares, instrument by instrument, with the public alternatives: faster and cheaper, certainly; accountable and impartial, structurally doubtfully - the regulator is a party to the matters it adjudicates, profits from the rules it writes, and faces, in place of elections, only the exit constraint that lock-in (Section 3.3) attenuates.

Governance as competition, and its public limits

Governance design is also competitive strategy, and the synthesis knits this chapter together. Platforms compete on their institutions: the marketplace with better fraud control, the app store with more trusted review, the social network with more predictable moderation is offering, precisely, superior law - and switching platforms is, among other things, choosing a legal system. This is the deep sense of Sussan and Acs's (2017) placeless ecosystems from Chapter 4 and of the platform society of van Dijck et al. (2018): the platform's rules are the business environment for its complementors and the civic environment for its users. Two limits define the frontier with Chapter 10. First, *conflicts of interest* are constitutive, not incidental: the regulator competes with the regulated (Zhu & Liu, 2018), taxes them (Debate 5.1), and adjusts the rules with both roles in hand - a fusion of powers no public constitutional tradition would tolerate, and the specific target of the DMA's self-preferencing and data-separation obligations. Second, *systemic spillovers* escape the platform's own optimization: the information disorder of Chapter 9 and the labour conditions of Chapter 7 are costs partly external even to the platform's long-run interest, which is why the public regulator's claim - to govern the governors - rests not on the platforms' failure to regulate, but on the structural mismatch between what they optimize and what a polity must. Discussion questions: (1) Take one platform's published policy document and classify each rule by the market failure it

addresses; which rules address no failure but a strategic interest, and how can an outsider tell? (2) Design the appeals process for seller suspension that you would mandate; cost it, and identify the caseload at which your design breaks. (3) ‘Choosing a platform is choosing a legal system.’ Develop the analogy - and locate exactly where it fails.

3.8. In Depth: The Economics of Interoperability

Debate 3.1 nominated interoperability as the remedy of choice for platform power; this section supplies the economics the debate presupposed, because interoperability is not one policy but a family, and its members have sharply different incentive structures.

Compatibility as strategy: the classical results

The founding analysis is Farrell and Saloner (1985): when network effects reward coordination on a common standard, decentralized adoption can fail in both directions - *excess inertia*, where users who would all gain from a new standard fail to switch because no one moves first (the coordination face of Section 3.3’s lock-in), and *excess momentum*, where a bandwagon stampedes to a new standard that leaves some adopters worse off. Communication and institutions - standards bodies, committees, announcements - mitigate both, which is why the digital economy is thick with the consortia and standards wars of Case 3.2. The strategic layer follows: compatibility is a *choice variable*, and the classical results predict who chooses what. Large incumbents favour incompatibility - their installed base is their moat, and interoperability would let entrants free-ride on network effects the incumbent built; small firms and entrants favour compatibility for the mirror reason; and the same firm flips sides as its position changes, a prediction the industry’s history confirms with almost comic regularity. Interoperability mandates, in this frame, are not neutral plumbing: they transfer the value of the installed base from its owner to the market, which is exactly why they are fought as property questions in regulatory proceedings.

The family of interoperabilities, and what each costs

Policy discussion gains precision by distinguishing at least three members of the family. *Data portability* - the right to extract one’s data and take it elsewhere (GDPR’s version, Section 10.2) - attacks individual switching costs but not the coordination problem: the user leaves with her history and without her network, so portability alone rarely moves markets where direct network effects dominate. *Protocol interoperability* - rival services interconnecting so users on one can transact or message with users on another, the telecommunications precedent and the

DMA's messaging obligation - attacks the network effect itself, converting competition *for* the network into competition *in* service quality; its costs are the ones Debate 3.1 rehearsed - a standing standards process, security surfaces across trust boundaries, and innovation pinned to the negotiated protocol's pace. *Adversarial interoperability* - entrants building compatibility without the incumbent's consent, by reverse engineering and scraping - supplied much of computing history's actual contestability (compatible hardware, alternative clients, migration tools) and has been progressively foreclosed by the legal perimeter of terms-of-service, anti-circumvention, and computer-misuse law: a reminder, in Lessig's (1999) terms, that the *legality* of interoperability is itself a regulatory choice, currently set, in most jurisdictions, in the incumbent's favour. The mature policy question is therefore compositional - which member of the family, applied at which layer, governed by whom - and the book's mechanisms supply the selection rule: match the instrument to the binding mechanism (portability for switching costs, protocols for direct network effects, adversarial rights for ecosystems too fast-moving for standardization), and budget honestly for each instrument's characteristic failure. Discussion questions: (1) Use Farrell and Saloner (1985) to explain why a dominant messaging platform resists interconnection that a challenger demands - and predict when the dominant firm's position will flip. (2) Why did number portability transform telecommunications while data portability has moved social networking little? Identify the mechanism each targets. (3) Draft the adversarial-interoperability safe harbour you would legislate; specify the security and privacy conduct it must still prohibit.

Chapter 4. Digital Entrepreneurship: Theory, Actors, and Ecosystems

Learning objectives. After studying this chapter, the reader should be able to: (1) situate digital entrepreneurship within the classical theories of the entrepreneur - Schumpeterian innovation, Kirznerian alertness, Knightian uncertainty - and explain what digital affordances add to each; (2) analyse the lean startup as a methodology of cheap experimentation and relate it to the economics of falling experimentation costs; (3) explain why digital ventures scale and internationalize at speeds unavailable to earlier firms, and what 'blitzscaling' owes to the platform economics of Chapter 3; (4) present the theory of entrepreneurial ecosystems and the role of venture capital in financing digital experimentation; (5) analyse the social, cultural, and institutional conditions of digital venturing, including the stratified access to entrepreneurship and the media construction of the founder; and (6) distinguish the venture-backed startup from the far larger population of platform-dependent micro-entrepreneurs.

KEY CONCEPTS

entrepreneurial opportunity; effectuation and experimentation; customer discovery; MVP; product-market fit; scalability; born-global venture; ecosystem; venture capital; runway

The platform economics of Chapter 3 described the structures of the digital economy; this chapter turns to its agents. Who builds the new firms, under what theory of action, with whose money, in which places, and against which social odds? The chapter braids the book's three lenses deliberately. Economics supplies the theory of the entrepreneurial function and of venture finance as a response to extreme uncertainty; sociology supplies the analysis of networks, regions, and stratified access; media studies supplies the critique of the founder as a cultural production - for few figures of contemporary capitalism are as intensively mediated as the startup entrepreneur.

4.1. From Schumpeter to Digital Affordances

Modern entrepreneurship research organizes itself around a triangle of classical positions. For Schumpeter (1934), as Chapter 1 noted, the entrepreneur is the innovator: the agent who carries out new combinations and thereby disrupts the circular flow of routine economic life, with profit as the temporary reward of successful disruption. For Kirzner (1973), the entrepreneur is the *discoverer*: markets are perpetually strewn with unnoticed profit opportunities - price discrepancies, unserved wants - and entrepreneurial alertness is the force

that finds them and, in finding them, drives markets toward coordination rather than away from it. For Knight (1921), the entrepreneur is the *bearer of uncertainty*: where risk can be calculated and insured, no entrepreneurial profit is possible; profit is the residual return to judgment exercised under genuine, unmeasurable uncertainty. Shane and Venkataraman (2000), in the field's canonical programmatic statement, synthesized these strands by defining entrepreneurship research as the study of the sources of opportunities and of the processes by which particular individuals discover, evaluate, and exploit them - the 'individual-opportunity nexus' that still frames the discipline.

Digitalization bears on every corner of the triangle. It is Schumpeterian in the most literal sense: each of the five waves of Chapter 1 unleashed a storm of creative destruction, and the recombinability of digital artefacts - software components, open-source libraries, public APIs, cloud services - makes 'new combinations' cheaper and faster to assemble than at any previous moment in economic history. It transforms Kirznerian discovery, because the data infrastructure of the digital economy renders opportunities newly visible: search trends, marketplace gaps, and user complaints are now legible at scale to anyone who looks. And it modulates, without abolishing, Knightian uncertainty: as the next section argues, digital tools convert some unmeasurable uncertainty into cheap testable hypotheses, even as the winner-take-most dynamics of Chapter 3 make the *distribution* of outcomes more extreme - a lottery with better diagnostics and bigger jackpots. To these classical bearings the contemporary literature adds the vocabulary of *affordances*. Nambisan (2017) and Autio et al. (2018), introduced in Chapter 1, argue that digital technologies do not merely lower the costs of entrepreneurship but alter its ontology: ventures remain permanently editable, boundaries between idea, product, and firm blur, and entrepreneurial agency disperses across ecosystems. Von Briel et al. (2018) systematize the point by treating specific digital technologies as *external enablers* whose properties - compression of time, conservation of resources, expansion of reach - can be mapped onto specific venture-creation processes, giving the affordance argument analytical teeth.

One classical insight deserves special emphasis because the celebratory literature so often forgets it. Baumol (1990) argued that the supply of entrepreneurial talent is roughly constant across societies; what varies is its *allocation* among productive, unproductive, and destructive uses, and the allocation is governed by the institutional structure of rewards. The digital economy furnishes Baumol's thesis with fresh illustrations on every side: the same alertness that builds a logistics platform can build a dark-pattern subscription trap, a click-farm, or a ransomware operation, and the same venture capital that funds productive experimentation

has funded regulatory arbitrage dressed as innovation. The question ‘how much entrepreneurship does the digital economy produce?’ is therefore always secondary to the question ‘toward what does its reward structure steer the entrepreneurship it produces?’ - a question that returns, with regulatory force, in Chapter 10.

COMMON MISCONCEPTION

An MVP is not the cheapest incomplete version of a product. It is the smallest credible experiment that can test a consequential assumption about a customer problem, value proposition, channel, willingness to pay, or usage behaviour.

4.2. Lean Startups, Scalability, and Born-Global Ventures

The managerial doctrine of digital venturing crystallized in the decade after the dot-com crash as the *lean startup*: a methodology of iterative experimentation in which the venture is treated not as a small version of a large firm but as a temporary organization searching for a viable business model (Blank, 2013; Ries, 2011). Its practical grammar - the minimum viable product, the build-measure-learn loop, validated learning, the pivot - institutionalizes a simple epistemology: founders hold hypotheses, not plans, and the cheapest possible test should be run before scarce capital is committed. Kerr et al. (2014) supply the economic frame: entrepreneurship is fundamentally *experimentation* under uncertainty, and the cost of running experiments is a first-order determinant of how much innovation an economy attempts. Digitalization collapsed that cost - cloud infrastructure replaced purchased servers, open-source components replaced bespoke code, online advertising replaced sales forces, app stores replaced distribution deals - and the lean method is best understood as the organizational form that cheap experimentation makes rational. The affinity with Knight is exact: the lean startup is a machine for converting portions of unmeasurable uncertainty into measured risk, one falsifiable hypothesis at a time.

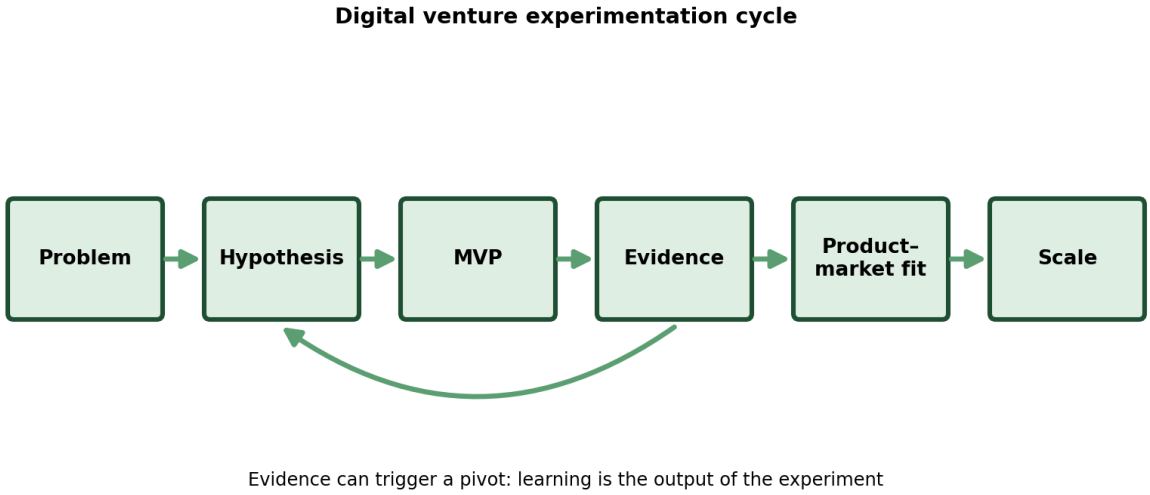

Figure 4.1. Digital venture experimentation: hypothesis, MVP, evidence, pivot and scale.

Cheap experimentation explains how digital ventures start; the economics of Chapters 2 and 3 explain how they grow. A firm whose product has near-zero marginal cost, whose distribution is the Internet, and whose market may tip to the leader faces a growth calculus unlike anything in the industrial economy: the constraint on expansion is neither factory capacity nor working capital for inventory but the speed at which users can be acquired before rivals acquire them. Hence the strategy the industry calls *blitzscaling (Hoffman & Yeh, 2018)* - the deliberate prioritization of speed over efficiency in pursuit of winner-take-most positions - with its signature pathologies of subsidized pricing, deferred profitability, and capital consumption on a scale that only the venture finance of Section 4.3 can sustain. The same cost structure dissolves the traditional geography of firm growth. Oviatt and McDougall (1994) had already identified, before the commercial web, the *international new venture* or born-global firm - organizations international from inception rather than after decades of domestic maturation; digital delivery made the type general. A software firm in Tallinn, Warsaw, or Bangalore is global on the day its product ships, a fact of considerable importance for the economic geography of Section 4.3 and for the development debates of Chapter 10.

Two caveats discipline the celebratory reading. First, scalability is a property of *some* digital businesses, not of digital business as such: ventures whose service embeds local physical operations - delivery, mobility, care - scale city by costly city, and many celebrated blitzscalers of the 2010s discovered that subsidized growth had purchased market share without purchasing a business model. Second, the lean method's epistemology has limits: iteration optimizes toward what measurable user behaviour rewards, which is not always what is socially valuable - a bias toward the engineering of engagement whose consequences for media and attention are

the subject of Chapter 9. Method, like talent, is allocatively neutral; Baumol's question applies to it with full force.

Applied Toolkit: From Opportunity to Product-Market Fit

A rigorous introduction to digital entrepreneurship should equip the reader not only to interpret ventures but also to interrogate their evidence. The practical unit of analysis is therefore the venture hypothesis: a falsifiable claim about a customer, a problem, a proposed value proposition, a channel, and an economically sustainable way of serving the customer. Customer discovery tests whether the problem is sufficiently frequent and costly to motivate action; solution testing asks whether the proposed intervention changes behaviour; product-market fit is reached only when adoption and retention cease to depend primarily on founder effort and become repeatable properties of the product-market relationship.

The useful sequence is problem → hypothesis → minimum viable product → observed behaviour → revision or pivot → repeatable acquisition and retention → scale. The minimum viable product is not the cheapest incomplete product but the smallest intervention capable of generating decision-relevant evidence. A good experiment therefore specifies in advance what observation would support the hypothesis, what observation would falsify it, and which decision follows from either result. This converts the lean startup from a slogan into an empirical method under uncertainty.

Market sizing should likewise be treated as a test rather than a decorative slide. Total addressable market (TAM) estimates the broad theoretical demand; serviceable available market (SAM) removes customers the venture cannot reach with its present product, geography, regulation, or distribution; serviceable obtainable market (SOM) is the share plausibly capturable under explicit assumptions about price, acquisition capacity and competition. Bottom-up estimates based on identifiable customers and purchasing frequency are usually more informative than top-down percentages of a large industry total.

Metric	Definition	Why it matters
CAC	Total acquisition spend / new customers acquired	Shows the price of growth and whether a channel scales.
LTV	Expected contribution margin over the customer relationship	Sets the economic ceiling for sustainable acquisition spending.
Churn	Share of customers or revenue lost in a period	Determines whether growth accumulates or merely replaces departures.

Retention cohort	Behaviour of users acquired in the same period	Separates product quality from changes in acquisition mix.
Burn rate	Net cash consumed per month	Translates strategy into time remaining.
Runway	Cash balance / monthly net burn	Defines the time available to reach the next financing or profitability milestone.

Table 4.1. Core venture metrics: the bridge from entrepreneurial narrative to testable economics.

4.3. Entrepreneurial Ecosystems and Venture Finance

Digital ventures do not germinate uniformly across space; they cluster, and the clusters persist. The contemporary vocabulary for this fact is the *entrepreneurial ecosystem*: the configuration of interdependent actors and factors - founders, investors, universities, skilled labour, service providers, exit markets, supportive culture - whose interaction in a territory enables productive venturing (Stam, 2015; Spigel, 2017). Spigel (2017) gives the concept its most careful formulation, analysing ecosystems as relational structures in which cultural attributes (stories of successful founders, tolerance of failure), social attributes (networks of mentorship, investment capital, worker mobility), and material attributes (universities, infrastructure, policy) reproduce one another over time. The intellectual pedigree runs through Saxenian's (1994) classic comparison of Silicon Valley and Boston's Route 128: two regions with equivalent endowments of technology, universities, and capital diverged because their industrial cultures differed - the Valley's dense, porous networks of inter-firm mobility and information exchange out-adapted Route 128's closed, vertically integrated corporations. The lesson, sociological to its core, is that the decisive infrastructure of innovation is relational: knowledge travels in people, and regions that let people circulate learn faster. Sussan and Acs (2017) extend the framework to the *digital* entrepreneurial ecosystem, in which the relevant habitat is only partly territorial - platforms, app stores, and open-source communities now function as placeless ecosystems whose governance (Chapter 3) conditions entrepreneurship as powerfully as any regional culture. Stam (2015) adds the necessary policy scepticism: 'ecosystem' passed into the strategy documents of every city and ministry faster than its causal claims were tested, and imitative attempts to transplant the Valley by building science parks have a long record of expensive failure.

The financial organ of the digital ecosystem is venture capital: intermediated risk capital invested in stages, in portfolios, against equity, in ventures too uncertain, too intangible, and too unprofitable for banks to touch. Gompers and Lerner (2001) describe the institutional innovations that make the form work - staged financing that preserves the option to abandon,

syndication that shares risk and pools judgment, convertible securities and board seats that align control with information, and the limited partnership that recycles institutional wealth through ten-year funds. The economics of experimentation (Kerr et al., 2014) explains the fit with digital venturing precisely: when experiments are cheap, failure is informative and inexpensive, and returns follow the extreme skew of winner-take-most markets - a few investments returning the fund, most returning little - a financier optimized for cheap abandonment and unbounded upside is the natural counterpart. But the fit is also a filter. Venture capital selects for the scalable, the tippable, and the exit-able; ventures that are merely good businesses - profitable, durable, local - fall outside its lens, as do, systematically, founders outside its networks. The geography of venture capital is more concentrated than the geography of entrepreneurship itself, and its network character - deals flow through referral, and referral flows through existing ties - makes it a powerful transmitter of the social stratification examined in the next section. Kenney and Zysman (2016) add a structural observation that binds this chapter to the last: in the platform era, venture-financed firms increasingly compete *for* intermediary positions, so that the capital burned in blitzscaling races is, in effect, the price of contesting a private regulatory throne.

4.4. Social, Cultural, and Institutional Conditions of Digital Venturing

Who gets to be a digital entrepreneur? The mythology of the open laptop - anyone, anywhere, can found the next great firm - collides with a stratification literature whose findings are unambiguous: entrepreneurship in general, and venture-backed digital entrepreneurship in particular, is patterned by class, gender, ethnicity, and network position. The mechanisms are cumulative rather than singular. Founding requires runway - savings or family wealth to survive unpaid experimentation - which screens by class before any investor is met. Granovetter's (1985) embeddedness, introduced in Chapter 1 as a general truth about markets, operates here with particular force: opportunities, co-founders, early customers, and above all capital flow through social networks, and the referral-driven deal flow of venture capital reproduces the demography of existing networks with mechanical reliability. The result, visible in every survey of venture funding, is that capital flows overwhelmingly to founders who resemble, in gender, ethnicity, and educational pedigree, the investors who fund them - not necessarily through overt animus, but through the ordinary sociology of homophily operating at the point where judgment under uncertainty must lean on familiarity. Baumol's allocation question thus acquires a distributional twin: institutions determine not only toward what entrepreneurship is steered, but from whom it is drawn - and talent excluded by the filter is simply lost.

Gender offers the best-documented example of this filter. Women found a minority of venture-backed firms, and the gap is not explained by the quality of their ventures alone. Using private data on interactions between investors and fundraising startups on an online investment platform, Ewens and Townsend (2020) find that male investors express less interest in female-led startups than in observably similar male-led ones, whereas female investors show the opposite pattern. Moreover, the male-led startups that male investors favour go on to underperform the female-led startups they favour. The evidence is therefore consistent with bias operating through homophily rather than with efficient screening, and it has an economic cost as well as a distributional one: ventures that are never funded because their founders do not resemble the investors are invisible losses of the ecosystem.

Culture completes the institutional picture, and here the media lens becomes indispensable. The founder is not merely an economic role but a mediated persona - produced by the business press, the conference stage, the biopic, and the platforms themselves - and the production has consequences. Marwick (2013), in an ethnography of the Silicon Valley technology scene, documents how social media transformed entrepreneurial self-presentation into continuous personal branding: status in the scene accrues to visibility, visibility demands the performance of authenticity, and the boundary between building a company and performing the building dissolves. The mediated founder-myth does economic work: it recruits talent into the startup labour market on lottery terms, legitimates the deferral of wages for equity, and frames regulatory forbearance as deference to genius. It also travels: policy discourses from Berlin to Bangalore import the Valley narrative wholesale, often detached from the institutional conditions - deep capital markets, permissive bankruptcy, the non-enforceability of non-compete clauses - that made the original possible. A sober institutional economics and a critical media studies here reach the same conclusion by different routes: the entrepreneurial culture that ecosystems require cannot be conjured by narrative alone, and narrative unmoored from institutions produces cargo-cult ecosystems and precarious founders.

The chapter closes by widening the frame. The venture-backed startup, for all its analytical and cultural prominence, is a demographic sliver of digital entrepreneurship. The numerically dominant form is the *platform-dependent entrepreneur* introduced in Chapter 1 and structurally analysed in Chapter 3: the marketplace seller, the app developer, the ride-hail driver-owner, the content creator - formally independent proprietors whose market access, pricing environment, and effective regulation are set by the private governance of platforms. For this population, the classical questions of entrepreneurship theory - opportunity, judgment, uncertainty - remain live, but they are answered within rules the entrepreneur neither makes

nor, often, sees; and the line between entrepreneurship and dependent labour, so clear in the textbooks, becomes the contested terrain on which the labour analyses of Chapter 7 and the regulatory battles of Chapter 10 are fought.

WORKED EXAMPLE

Scenario. A startup has EUR 300,000 cash. Monthly operating costs are EUR 52,000 and recurring gross-margin contribution is EUR 17,000.

Calculation. Net burn = EUR 35,000 per month. Runway = 300,000 / 35,000 = approximately 8.6 months.

Decision. If the next financing round normally requires six months of preparation, the firm has little room for an unfocused feature roadmap. Its experiments should target the assumptions most likely to change retention, willingness to pay, or acquisition efficiency.

THREE LENSES

Economics asks: How is uncertainty converted into option value through staged experiments and finance?

Sociology asks: How do networks, status, geography, and institutional access shape who can become a founder?

Media studies asks: How do founder narratives, visibility, and platform-mediated reputation influence capital and demand?

4.5. Chapter Summary and Discussion Questions

This chapter has assembled an interdisciplinary account of digital entrepreneurship. The classical triangle - Schumpeterian innovation, Kirznerian discovery, Knightian uncertainty, synthesized in Shane and Venkataraman's (2000) opportunity nexus - was shown to be transformed rather than replaced by digital affordances (Nambisan, 2017; von Briel et al., 2018), with Baumol's (1990) allocation thesis supplying the critical counterweight. The lean startup was interpreted through the economics of cheap experimentation (Kerr et al., 2014), and the scalability and born-global character of digital ventures (Oviatt & McDougall, 1994) were derived from the cost structures and tipping dynamics of Chapters 2 and 3. Entrepreneurial ecosystems were analysed relationally (Saxenian, 1994; Spigel, 2017; Sussan & Acs, 2017), venture capital was presented as the financial technology of skewed experimentation (Gompers & Lerner, 2001), and the social and cultural conditions of venturing - stratified access, networked capital, and the mediated founder-myth (Marwick, 2013) - completed the picture, before the frame was widened to the platform-dependent entrepreneurs who constitute the phenomenon's demographic majority.

Discussion questions. (1) Apply the Schumpeter-Kirzner-Knight triangle to a digital venture you know: which theory best describes its founding moment, and what does each miss?

(2) ‘The lean startup is Knightian uncertainty domesticated by cheap experiments.’ Explicate and evaluate. (3) Why does venture capital fund blitzscaling? Connect the skew of venture returns to the tipping conditions of Chapter 3, and identify a market where the strategy destroyed value. (4) Using Spigel (2017) and Saxenian (1994), diagnose your own city or region as an entrepreneurial ecosystem: which attributes are present, which absent, and which could policy realistically supply? (5) Venture funding flows disproportionately to founders demographically similar to investors. Distinguish taste-based, statistical, and network-structural explanations, and propose an intervention targeting each. (6) Is the marketplace seller an entrepreneur? Defend both answers, specifying the theoretical criterion each relies on, and state what turns on the classification.

4.5a. In Focus: Digital Entrepreneurship beyond Silicon Valley

The canonical cases in this chapter are drawn largely from the United States and Europe, yet some of the most instructive digital ventures of the last two decades emerged in the Global South, often where weak legacy infrastructure made mobile-first "leapfrogging" possible. M-Pesa (Kenya, from 2007; see Case 4.1) turned a mobile-network operator into a payments rail and brought basic financial services to millions who had never held a bank account. India's Unified Payments Interface (UPI), together with the Jio mobile rollout, reshaped a market of a billion users through public digital infrastructure and aggressive data pricing. Nubank (Brazil) built one of the world's largest digital banks by attacking a concentrated, high-fee incumbency, and Gojek and Grab (Indonesia and Southeast Asia) bundled ride-hailing, payments and dozens of on-demand services into "super-apps" with no close Western analogue.

These cases show that the affordances discussed in Section 4.1 combine differently with local institutions: the binding constraint is often trust, distribution or regulation rather than technology, and the successful venture is frequently the one that supplies a missing institution rather than assuming it exists. Seminar question: choose one venture above, identify the digital affordance it exploited, and name the absent institution it had to build itself.

4.5b. Numerical and Data Exercises

Exercise 4.A (power-law returns). A seed fund invests EUR 1 million in each of 20 startups. After ten years, 12 are written off, 5 return their invested capital (1×), 2 return 3×, and one returns an unknown multiple X. (a) What must X be for the fund to return 3× its capital (EUR 60 million) before fees? (b) What share of total proceeds does the best investment then contribute? (c) Relate the answer to the fund-returner criterion of Section 4.7.

Exercise 4.B (dilution and runway). Founders own 100% of a startup. A seed round raises EUR 2 million at a pre-money valuation of EUR 8 million; a Series A later raises EUR 5 million at a pre-money valuation of EUR 20 million. (a) What are the founders' and the seed investors' stakes after each round? (b) If monthly net burn after the Series A is EUR 250,000 and the company holds EUR 5.5 million, what is its runway? (c) Why does staged financing give investors an option to abandon (Section 4.3)?

Indicative answers. 4.A (a) The other investments return $5 \times 1 + 2 \times 3 =$ EUR 11 million, so $X = 60 - 11 = 49$, i.e. 49×. (b) $49 \div 60 \approx 82\%$ of all proceeds. (c) A fund can reach its target only if some investment can return the whole fund; investments that cannot, however safe, do not fit the model. 4.B (a) Seed: post-money EUR 10 million, so investors hold 20% and

founders 80%. Series A: post-money EUR 25 million, so new investors hold 20%, founders 80% × 0.8 = 64% and seed investors 20% × 0.8 = 16%. (b) 5.5 ÷ 0.25 = 22 months. (c) Each round is a new decision taken with more information; investors can decline to continue if milestones are missed, which limits losses.

4.6. Case Studies and Debates

The materials below test the entrepreneurship toolkit against three settings the Silicon Valley template does not describe: a financial innovation from Nairobi, a blitzscaling failure from the venture heartland itself, and the standing controversy over the state's role in the innovation system.

Case 4.1. M-Pesa: digital entrepreneurship without the template

The most consequential digital-financial innovation of the smartphone era's eve arrived from none of the celebrated ecosystems. M-Pesa, launched in Kenya in 2007 by the telecommunications operator Safaricom, converted the mobile phone - then overwhelmingly a basic handset - into an instrument for storing and transferring money through a network of human agents who exchange cash for electronic value. Adoption was explosive by any standard in the diffusion literature: within a few years the service reached the large majority of Kenyan households, and mobile money became infrastructure - the rail on which wages, remittances, school fees, and eventually credit and savings products travel. Jack and Suri (2014) supply the canonical economic evaluation: by collapsing the transaction costs of transfers across distance, mobile money transformed risk-sharing - households with access absorbed income shocks through wider, faster remittance networks that unconnected households could not tap - a welfare effect operating through exactly the mechanism this book has emphasized since Chapter 2, the reduction of transaction costs, here applied to the informal insurance networks that embeddedness theory (Chapter 8) describes.

Read against Chapters 3 and 4, the case is a compendium. The launch solved a two-sided ignition problem - senders need receiving agents, agents need transaction volume - through a strategy the theory recommends: the operator leveraged an existing network (its airtime-reseller base became the agent network) and a marquee use case ('send money home') with urban-to-rural remittances as the beachhead. The innovation was institutional as much as technical: a permissive regulator allowed a telecommunications firm to perform bank-like functions ahead of banking law, illustrating Baumol's (1990) thesis from its productive side - the reward structure, briefly open, steered entrepreneurship into a transformative channel. And

the innovator was a corporate incumbent, not a startup, warning against equating entrepreneurship with the venture-backed form of Section 4.3. Seminar questions: (1) Why did mobile money ignite in Kenya and stall, for years, in markets with richer banking systems? Use the tipping conditions and the concept of institutional voids. (2) Was M-Pesa Schumpeterian, Kirznerian, or Knightian - and does the corporate identity of the innovator strain the theories? (3) M-Pesa's success made Safaricom a gatekeeper of Kenyan payments; apply Chapter 3's complementor analysis and Chapter 10's regulatory options.

Case 4.2. WeWork and the limits of blitzscaling

If M-Pesa shows the toolkit travelling well, the collapse of WeWork's 2019 public offering shows what happens when the toolkit is applied to a business it does not fit. The firm leased office space long-term, refurbished it, and subleased it short-term with services attached - an arbitrage on duration and density with a century of low-margin precedent. Its venture-era narrative, however, was platform-shaped: the company presented itself as a technology firm with network effects, community as moat, and winner-take-most geography, and it raised and consumed capital on that theory at blitzscaling rates, reaching a private valuation in the tens of billions before public-market scrutiny of its prospectus - persistent losses scaling with revenue, no visible network effect, governance concentrated in the founder - collapsed the offering, the valuation, and eventually the founder's tenure within weeks. The company reached public markets only in 2021, through a merger with a special-purpose acquisition company, and filed for bankruptcy protection in 2023.

The analytical value of the episode is diagnostic. Chapters 2 and 3 specify when blitzscaling is rational: near-zero marginal costs, demand-side increasing returns, and a market that tips. Office subleasing has none of these - each new desk costs what desks cost, a tenant's value does not rise with other tenants in another city, and nothing prevents multi-homing among landlords - so capital spent on speed purchased scale without purchasing defensibility: the strategy's cost with none of its prize. The deeper lessons implicate the system rather than the firm: the abundance of late-cycle venture capital (Section 4.3) rewarded narrative isomorphism - dressing any business as a platform - and the mediated founder-myth of Section 4.4 supplied the costume. Kenney and Zysman's (2016) observation that venture capital competes to fund intermediary positions explains the temptation; the episode shows its failure mode. Seminar questions: (1) Write the two-page memo a sceptical analyst should have written in 2018, using only concepts from Chapters 2-4. (2) Distinguish this failure from that of a genuine platform that simply lost its tipping race; which observable pre-collapse indicators separate the cases?

(3) What changes in venture governance - board structure, disclosure, dual-class shares - would have surfaced the problem earlier, and at what cost to legitimate experimentation (Kerr et al., 2014)?

Debate 4.1. The entrepreneurial state

Motion: the state, not private venture capital, is the digital economy's foundational venture capitalist, and innovation policy should be designed - and its returns shared - accordingly. The affirmative case is Mazzucato's (2013): the general purpose technologies beneath every chapter of this book - the Internet's protocols, the microprocessor's early markets, satellite positioning, touchscreens, the algorithms behind search, the foundational research beneath modern AI - were funded at their riskiest stages by public agencies with mission-driven patience that ten-year funds (Section 4.3) structurally lack; private capital entered at deployment, captured the returns, and financed a mythology (Section 4.4) that erases its own subsidy. Justice and future capacity alike therefore demand that the public investor take equity-like returns - golden shares, royalties, conditionality on reinvestment - rather than socializing risk and privatizing reward. *Against the motion:* the state's successes are read with survivor bias from a portfolio whose failures were equally public and vastly numerous. Mission agencies fund technologies, but the market discovery of what users want - the Kirznerian and lean-experimental work of Sections 4.1 and 4.2 - is precisely what bureaucracies do worst. Attaching returns and conditions to public research would tax diffusion, the very externality public funding exists to subsidize. And the observable counterfactual is unkind - economies where the state selected digital champions directly have produced fewer frontier firms than those where it funded science and let entry sort the rest. The debate should force precision on the verb: *funding research*, *procuring missions*, *subsidizing deployment*, and *picking firms* are four different policies, and each side must state which it defends or attacks.

4.7. In Depth: Venture Finance Through the Cycle

Section 4.3 introduced venture capital as the financial technology of skewed experimentation; this advanced section opens the machinery. Its three movements - the contract, the portfolio, and the cycle - correspond to the three levels at which venture finance shapes digital entrepreneurship: the individual deal, the fund, and the macro-financial environment whose swings the reader met in Case 1.2 and Case 4.2.

The contract: allocating control under extreme uncertainty

The venture deal is a solution to a contracting problem that would defeat ordinary finance: the entrepreneur knows more than the investor, the asset is intangible and inseparable from its founders, the outcome distribution is wildly skewed, and the venture will need repeated infusions whose terms cannot be fixed in advance. Kaplan and Strömberg (2003), in the canonical empirical study of actual venture contracts, documented how practice answers theory: venture agreements systematically *unbundle* the rights that ordinary equity bundles together - cash-flow rights, voting rights, board seats, and liquidation preferences are allocated separately, and, crucially, *contingently*: control shifts with performance, so that founders who hit milestones retain autonomy while underperformance transfers board control to investors. Convertible preferred stock is the workhorse instrument because it encodes the contingency - downside protection through the preference and the board, upside participation through conversion - and staged financing (Section 4.3) supplies the enforcement: each round is an option to abandon, so the entrepreneur is disciplined not by covenants but by the perpetual need to earn the next tranche. The design maps precisely onto the theory of Section 4.1: Knightian uncertainty cannot be priced, so the contract prices *governance* instead, allocating decision rights to whoever the state of the world reveals to be the better judge. The founder mythology of Section 4.4 rarely mentions that its heroes typically operate inside this cage of contingencies - and the celebrated exceptions, the founders who retained control through dual-class shares, are exceptions purchased with bargaining power that only the hottest deals confer, with Case 4.2's governance collapse as the cautionary countercase.

The portfolio: power laws and their behavioural consequences

The fund-level arithmetic drives behaviour that appears pathological until the distribution is understood. Venture returns are not merely skewed but power-law distributed: in a typical successful fund, a majority of investments lose money, a modest group returns capital, and one or two investments generate most of the fund's entire return. Three consequences

follow with mechanical force. First, the *fund-returner criterion*: an investment that can succeed only moderately is, for the venture investor, equivalent to a failure - the fund's economics require each deal to be capable, however improbably, of returning the fund - which is why solid, profitable, non-scalable businesses (Section 4.2's caveat) are declined not despite but because of their moderate reliability, and why founders are pressed toward the blitzscaling gamble even when a smaller, surer business was available. Second, *failure tolerance as policy*: because the option to abandon is the portfolio's risk management (Kerr et al., 2014), the ecosystem develops the institutions of cheap failure - acqui-hires, soft landings, the destigmatized shutdown - that Section 4.3's ecosystem literature treats as cultural attributes and this section derives from arithmetic. Third, *herding on the tail*: since returns concentrate in outliers, and outliers are identified partly by social proof, the deal flow homophily of Section 4.4 has a financial accelerant - missing the outlier is catastrophic, backing a consensus failure is forgivable - producing the thematic stampedes (into whatever category last produced a giant) that each cycle exhibits.

The cycle: cheap money, dry powder, and the discipline of exits

Venture finance breathes with the macro-financial cycle, and the digital economy's history is legible in its lungs. The fund structure transmits the cycle: institutional investors allocate to venture as a risk asset, so low interest rates and buoyant public markets swell fund-raising. Funds must deploy within their investment periods, so swollen fund-raising becomes 'dry powder' pressing valuations upward and diligence downward. Finally, the exit market - public offerings and acquisitions - converts paper marks into realized returns, so a closed exit window (a falling market, a hostile antitrust regime toward acquisitions, per Chapter 10) backs up the entire pipeline, stranding late-stage ventures priced for a world that ended. The 2010s ran the mechanism at historic scale: a decade of near-zero rates pushed unprecedented capital into venture and into its late-stage imitators - sovereign funds, crossover hedge funds, corporate vehicles - compressing the discipline of staged finance (mega-rounds replaced milestones) and financing the subsidy wars of Section 4.2; the rate reversal of the early 2020s ran it in reverse, with down rounds, shutdowns, and a rediscovery of unit economics. Perez's (2002) frame from Case 1.2 supplies the long view - finance periodically overshoots the technology it funds, and the overshoot builds infrastructure - while the AI investment wave contemporary with this book invites the reader to apply the diagnostic apparatus in real time rather than in retrospect. Discussion questions: (1) Using Kaplan and Strömberg (2003), explain why venture contracts allocate control contingently rather than proportionally to ownership; what would proportional

control misprice? (2) Derive the fund-returner criterion from a power-law return distribution, and use it to explain a real refusal: why will a venture fund not finance an excellent restaurant chain? (3) Trace one channel by which a central-bank rate decision reaches a seed-stage founder's term sheet. (4) 'Venture capital does not fund innovation; it funds the pursuit of monopoly.' Assess, drawing on Sections 3.3, 4.2, and this section.

4.8. In Depth: Cooperative and Social Alternatives - Owning the Platform Otherwise

This chapter has analysed digital entrepreneurship in its dominant, investor-owned form, and Section 4.3 showed how venture finance filters what gets built; this section examines the road not funded - platform cooperativism and social digital entrepreneurship - both because the alternatives illuminate, by contrast, the institutional specificity of the mainstream, and because their persistent difficulties are themselves a lesson in the book's mechanisms.

The cooperative proposition

The proposition, given its manifesto by Scholz (2016), is structurally simple: the platform's code is replicable (Chapter 2), so the extractive part of the platform economy is not the technology but the ownership - replace investor ownership with ownership by the workers, users, or producers whom the platform coordinates, and the surplus, the data, and the governance of Sections 3.7 and 7.1 accrue to those who generate them. The lineage is older than the terminology: agricultural marketing cooperatives, mutual insurers, and credit unions solved analogous intermediation problems under member ownership for a century, and the proposal reads the gig platform as the newest case of an old form - the labour-supplying members hiring the coordinator rather than the reverse. Scattered working examples exist across ride services, delivery collectives, freelancer platforms, stock-photo cooperatives, and community-owned social networks; their existence proves feasibility, and their scale proves something else, which is the section's analytical centre.

Why the alternative stays small: the mechanisms turned around

Read with this book's toolkit, the cooperative's obstacles are the mainstream's advantages inverted, and enumerating them is the best available stress test of the whole argument. *Capital*: the winner-take-most races of Section 4.2 are financed by venture capital's tolerance for years of losses (Section 4.7); cooperatives, structurally unable to sell control, cannot raise blitzscaling money, and so contest tipped markets against subsidized incumbents on member dues - network-effect competition (Chapter 3) with one side's war chest removed. *Ignition*: the chicken-and-egg problem (Section 3.2) is indifferent to virtue; the cooperative must still assemble both sides, and mission attracts members more reliably than it attracts the marginal customer choosing on price and liquidity. *Governance*: member democracy is the point and also a cost - heterogeneous member interests slow the pivots that Section 4.2's method demands - though the platform-cooperative rejoinder deserves equal weight: the investor-

owned platform's 'agility' includes the unilateral rule changes of Section 3.4, which is precisely the arbitrariness the cooperative form exists to remove. The plausible niches follow from the analysis rather than from sentiment: markets that are local rather than tippable (city-scale services where global network effects are weak), members with aligned interests and repeated dealings, and - the increasingly important case - *protocol-layer* alternatives, where open standards and federated architectures (Section 3.8's interoperability by design) substitute for the coordinating firm altogether, shifting the question from who owns the platform to whether a platform need be owned at all. The policy postscript writes itself from Chapter 10: if the filter is finance and the barrier is tipping, then cooperative-friendly capital instruments, interoperability mandates, and public procurement are the levers - and Baumol's (1990) thesis applies once more: the share of entrepreneurship that takes cooperative form is, like everything else in this book, an institutional variable, not a fact of nature. Discussion questions: (1) Take one gig market and design its cooperative competitor: capitalization, ignition strategy, governance. Identify the mechanism that kills it, and the regulation that would save it. (2) 'The cooperative's problem is not efficiency but finance.' Assess using Sections 4.3 and 4.7. (3) When does federation (protocols) dominate cooperation (ownership) as the alternative to the investor platform? Use Section 3.8.

Chapter 5. Business Models of the Digital Economy

Learning objectives. After studying this chapter, the reader should be able to: (1) define the business model concept and explain why digitalization made business-model innovation a distinct locus of competition; (2) analyse the advertising-funded, subscription, and freemium models as applications of the pricing economics of Chapter 2, and evaluate their behavioural and social externalities; (3) distinguish marketplaces from resellers and from aggregators, and explain the economics of on-demand and sharing-economy services; (4) describe the software-as-a-service and cloud models and their consequences for industry structure and entrepreneurship; (5) analyse the creator economy as a system of platformized cultural production; and (6) assess any digital business model by asking, systematically, who pays, who is subsidized, what is measured, and where the risk sits.

KEY CONCEPTS

business model; value creation and capture; advertising; subscription; freemium; marketplace take rate; SaaS; churn; CAC; LTV; contribution margin; creator economy

The preceding chapters supplied mechanisms: cost structures, network effects, price structures, experimentation. This chapter assembles them into the recurring configurations by which digital firms actually capture value - their business models. The concept deserves a moment of definition, since it entered management scholarship in earnest only with the digital economy itself. Teece (2010) defines a business model as the architecture by which a firm creates value for customers and converts a share of it into profit - the design of what is offered, to whom, how it is delivered, and how it is monetized. Zott et al. (2011), reviewing the literature's rapid growth, note that the concept became indispensable precisely when digitalization decoupled these design choices from one another: the same underlying technology can be monetized as an advertising business, a subscription, a marketplace, or a licence, and the choice among them - not the technology - frequently decides commercial fate. Practitioner frameworks such as Osterwalder and Pigneur's (2010) business model canvas institutionalized this insight as a design discipline, and the lean startup of Chapter 4 is, at bottom, a method for searching the space of business models experimentally. The taxonomy that follows is therefore not a list of industries but a grammar of monetization, and its recurring analytical questions are four: who pays, who is subsidized, what is measured, and who bears the risk.

5.1. Advertising-Funded, Subscription, and Freemium Models

The advertising-funded model is the oldest and, by user reach, still the dominant monetization of the consumer Internet: the service is free to users, and revenue is earned by selling advertisers access to the audience the free service assembles. Chapter 3 supplied its formal skeleton - a two-sided market with the user side fully subsidized - and media history supplies its pedigree: commercial broadcasting and the penny press ran the same structure, selling audiences to advertisers, a continuity on which Smythe's (1977) audience-commodity thesis and Wu's (2016) history of the 'attention merchants' both insist. What digitalization changed was the precision of the sale. Where broadcast advertisers bought demographic approximations, digital advertisers buy individually targeted impressions, priced in real-time auctions and attributed to subsequent behaviour - a transformation built directly on the collapse of tracking costs described in Chapter 2 (Evans, 2009). The economics are formidable: targeting raises the value of each impression, auction mechanisms extract that value efficiently, and measurement closes the loop that print and broadcast never could. So are the externalities. The model's input is attention and its raw material is behavioural data, which yields two structural pressures examined at length in later chapters: the engineering of engagement, since time-on-service is inventory (Chapter 9), and pervasive data collection, since targeting improves with surveillance (Chapter 6). Goldfarb and Tucker (2011) demonstrate the coupling econometrically from the reverse direction: when European privacy regulation restricted tracking, the measured effectiveness of display advertising fell substantially - evidence that the model's profitability and its data appetite are not separable.

The subscription model inverts the payment structure: users pay directly, typically a recurring flat fee for access to a service or catalogue. Its economics combine the bundling logic of Chapter 2 - the streaming catalogue is Bakos and Brynjolfsson's (1999) large-number bundle in commercial form - with the finance of recurring revenue: predictable cash flows, valuation by subscriber metrics such as lifetime value and churn, and competition refocused from the transaction to the relationship. Strategically, subscription realigns incentives in ways advertising cannot: a subscriber-funded service maximizes perceived ongoing value rather than engagement-time, which is why the model is repeatedly proposed as the remedy for advertising's pathologies - and why its limits matter. Subscription fatigue, the regressive character of paywalls, and the exclusion of those unable or unwilling to pay reintroduce, in commercial form, the access questions that zero-price services had seemed to dissolve; in journalism, the subscription turn has produced well-resourced titles behind paywalls and an advertising-starved remainder outside them, a stratification of the information commons taken up in Chapter 9.

Freemium, the third canonical configuration, is versioning (Chapter 2) deployed as an acquisition machine: a functional free tier recruits users at zero marginal cost, and conversion to paid tiers monetizes the minority with the highest willingness to pay. The model's arithmetic is unforgiving - conversion rates in low single digits are typical, so the free base must be enormous and its marginal cost near zero - which is why freemium is native to software and services and rare where the free tier consumes real resources. Its design problem is the versioning problem in acute form: the free tier must be good enough to spread and incomplete enough to convert, and the boundary is a strategic variable tuned as deliberately as any price. In games and consumer applications, freemium has evolved a behavioural variant - free products monetized through in-app purchases engineered around compulsion loops - whose consumer-protection and design-ethics controversies illustrate, once more, Baumol's allocation thesis applied to business-model talent.

COMMON MISCONCEPTION

Growth is not a business model. User growth can destroy value when marginal acquisition costs exceed expected contribution, when retention is weak, or when platform subsidies are masking negative unit economics.

5.2. Marketplaces, Aggregators, and On-Demand Services

The marketplace model monetizes intermediation itself: the platform matches independent buyers and sellers and takes a commission on transactions it enables. Chapter 3 established its structural economics - two-sidedness, ignition, governance - and the managerially decisive design choice it frames: whether to operate as a *marketplace*, in which third parties set prices and hold inventory risk, or as a *reseller*, which buys, prices, and sells on its own account. Hagiu and Wright (2015) formalize the trade-off as an allocation of control rights: marketplace mode economizes on the platform's information about local demand and taps sellers' incentives and variety; reseller mode secures the buyer experience and internalizes spillovers among products. Most large platforms are hybrids, and the mode boundary is itself strategic - the marketplace that observes its sellers' data and then enters their categories as a reseller is the envelopment problem of Chapter 3 restated as business-model design (Zhu & Liu, 2018).

The on-demand and 'sharing economy' services that spread with the smartphone - ride-hailing, short-term accommodation, delivery, freelance labour markets - are marketplaces whose supply side consists of distributed individuals and their underused assets or time. Einav et al. (2016) supply the economic anatomy: peer-to-peer markets exist where technology can

solve, cheaply and at scale, the three problems that once confined such exchange to acquaintances - search and matching, pricing, and trust - the last through the reputation mechanisms whose economics Chapter 2 introduced and whose sociology Chapter 8 examines. Sundararajan (2016) reads the phenomenon optimistically as 'crowd-based capitalism': the mobilization of idle capacity, lower entry barriers to micro-entrepreneurship, and consumption shifted from ownership to access. The critical literature replies that the sharing framing romanticizes what is, in revenue terms, mostly on-demand service labour intermediated by venture-financed firms; that the 'assets' shared soon professionalize into dedicated commercial stock; and that the model's unit economics frequently depend on reclassifying employment as independent contracting - the labour-law battleground of Chapter 7. Both readings are correct about different parts of the phenomenon, which is precisely why the business-model lens matters: 'the sharing economy' names not one model but a family, distinguishable by who holds the asset, who bears demand risk, and who sets the price.

A final intermediary type deserves separate notice because its power does not rest on transactions at all. The *aggregator* - the search engine, the app store front page, the social feed, the recommendation surface - monetizes neither goods nor commissions primarily, but the allocation of attention across others' offerings. Its business model is the auction or ranking of visibility, and its structural position - owning demand while suppliers compete for placement - gives it leverage over producers that resembles, in economic function, the shelf-space power of dominant retailers raised to informational completeness. The aggregator is where the business models of this chapter meet the media theory of Chapter 9: what it sells is reach into the public's finite attention, and how it ranks is therefore simultaneously a commercial algorithm and an editorial act.

5.3. Software-as-a-Service and Cloud Infrastructures

The quietest of the great business-model transitions was the migration of software from product to service. Packaged software - the second wave's characteristic commodity - was sold as a perpetual licence: a large upfront price, installation on the customer's machines, and upgrade revenue extracted release by release. Software-as-a-service (SaaS) replaces the licence with a subscription to an application operated by the vendor and delivered through the browser: revenue becomes recurring, deployment instantaneous, and every customer runs the current version because there is only the current version. The model's economics recombine familiar elements - near-zero marginal cost of serving an additional tenant, subscription finance with its churn-and-lifetime-value calculus, and versioned tiers from free trials to enterprise contracts -

but its strategic consequence is distinct: the vendor now observes usage continuously, so the build-measure-learn loop of Chapter 4 runs against live customers, and the product improves by telemetry. The corresponding dependency is equally distinct: the customer who stops paying loses not an old version but the software altogether, and frequently the switching costs of accumulated data and integrations besides - lock-in (Chapter 3) implemented as a billing model.

Beneath SaaS lies the deeper stratum: cloud computing, the provision of computing itself - processing, storage, networking, and increasingly machine-learning capacity - as a metered utility. Armbrust et al. (2010), in the reference account, identify the model's essence in three properties: the appearance of infinite computing resources on demand, the elimination of upfront commitment, and payment for use at fine granularity, all resting on the massive economies of scale of hyperscale data centres. The transformation this worked on entrepreneurship has been noted in Chapter 4 - capital expenditure became operating expenditure, and the cost of founding a software firm collapsed - but the industry-structure consequence deserves equal emphasis: computing infrastructure consolidated into a handful of global providers whose platforms sit beneath much of the digital economy, including many of their customers' competitors. The cloud provider is thus the ultimate enveloper-in-waiting of Chapter 3 - it observes workloads, supplies primitives, and ascends the stack into any service its scale makes attractive - and, as Chapter 10 will argue, a concentration of infrastructural power that competition policy designed for product markets struggles even to describe. The pattern completes a historical irony worth savouring: an industry born by decentralizing the mainframe has, two waves later, rebuilt the mainframe at planetary scale and rented it back.

5.4. The Creator Economy and Direct Monetization

The newest layer of the business-model stack monetizes neither software nor intermediation but persons: the creator economy, in which individuals produce media - video, audio, text, performance, instruction - distributed through platforms and monetized through a widening menu of instruments: advertising revenue shares, brand sponsorship, subscriptions and memberships, tips and virtual gifts, commerce and merchandise, and paid content. Cunningham and Craig (2019) name the phenomenon *social media entertainment* and document its scale: a globally distributed screen industry, adjacent to but organizationally unlike traditional media, in which previously amateur creators professionalize atop platform infrastructure. Analytically, the creator is the platform-dependent entrepreneur of Chapters 3 and 4 in cultural production: formally independent, algorithmically distributed, and governed by terms - revenue splits, monetization eligibility, recommendation criteria - that the platform

sets and revises unilaterally. Nieborg and Poell (2018) give this condition its theoretical name, the *platformization of cultural production*: cultural goods become contingent commodities, perpetually reshaped by platform metrics, and their producers become optimizers of signals they can observe but not control.

The sociology of creator work has been documented with particular force by Duffy (2017), whose study of aspirational labour describes the model's human economics: a vast base of creators works unpaid or underpaid in the hope of future monetization, sustained by the visible success of a tiny apex and by an ideology - do what you love - that reframes speculative labour as self-fulfilment. The income distribution of creator platforms is, predictably from Chapter 3, extremely skewed: winner-take-most dynamics operate on attention just as on markets, recommendation algorithms amplify the already-visible, and the median creator earns little while the distribution's tail earns spectacularly. Direct-monetization instruments - memberships, tips, paid newsletters - moderate the skew at the margin by converting a creator's most committed audience into recurring patrons, shortening the value chain from advertiser-mediated to audience-direct; they also re-expose the creator to subscription economics, churn included, and leave the platform's toll intact. The creator economy thus condenses the book's argument to date into a single figure: the person who is at once a media outlet (Chapter 9's concerns), a micro-enterprise (Chapter 4's), a complementor (Chapter 3's), and a worker (Chapter 7's) - and whose livelihood depends on pricing, ranking, and policy decisions made by the infrastructure on which she performs. Poell et al. (2022) offer the synthetic treatment, and Chapters 7 and 9 return to the labour and media faces of the figure respectively.

Applied Toolkit: Reading a Digital Business Model

A digital business model can be audited by separating four questions that are often conflated: who receives the value proposition, who pays, which variable the firm optimizes, and where residual risk is placed. The same technical service may support radically different institutional arrangements depending on those answers. Students should therefore resist labels such as 'platform', 'subscription' or 'creator economy' until the cash flow, control rights and risk allocation have been mapped explicitly.

Model	**Who pays?**	**Key operating metric**	**Characteristic risk**
Advertising-funded	Advertiser	Revenue per impression / user; incremental lift	Engagement pressure, surveillance, attribution error

Subscription	User or organization	Retention, churn, LTV	Subscription fatigue and access exclusion
Freemium	Minority of converting users	Free-to-paid conversion; marginal serving cost	Large non-paying base; incentive to degrade free tier
Marketplace	Transaction side(s)	Liquidity, take rate, successful matches	Chicken-and-egg launch; disintermediation
SaaS	Recurring customer	MRR/ARR, net revenue retention, CAC payback	Switching costs, concentration in cloud dependencies
Creator/direct	Audience, sponsors, platform	Audience concentration, conversion, platform take	Algorithmic dependency and volatile reach

Table 5.1. Comparative anatomy of recurring digital business models.

WORKED EXAMPLE

Scenario. A SaaS product charges EUR 50 per month, has 80% gross margin, monthly logo churn of 2%, and CAC of EUR 600.

Calculation. A simple steady-state LTV approximation is (50 × 0.80) / 0.02 = EUR 2,000. LTV/CAC = 3.33. CAC payback is approximately 600 / 40 = 15 months.

Interpretation. The ratio looks attractive, but the 15-month payback means cash requirements may still be severe during rapid growth. Unit economics and financing economics must be read together.

THREE LENSES

Economics asks: Who pays, for what, with which marginal cost and pricing mechanism?

Sociology asks: Which forms of trust, identity, and community sustain participation and payment?

Media studies asks: How does the model shape content incentives, visibility, creative autonomy, and audience relationships?

5.5. Chapter Summary and Discussion Questions

This chapter has organized the digital economy's recurring monetization architectures. The business-model concept was defined (Teece, 2010; Zott et al., 2011) and its rise linked to digitalization's decoupling of technology from monetization. The advertising-funded model was analysed as a subsidized two-sided structure with precision targeting built on cheap tracking (Evans, 2009), and its dependence on data documented econometrically (Goldfarb & Tucker, 2011); subscription was presented as bundling plus recurring-revenue finance, and freemium as versioning deployed for acquisition. Marketplace economics were developed through the marketplace-versus-reseller choice (Hagiu & Wright, 2015), the anatomy of peer-to-peer markets (Einav et al., 2016) and the contested 'sharing economy' framing (Sundararajan, 2016), and the aggregator was isolated as the intermediary that monetizes

attention allocation itself. SaaS and cloud computing (Armbrust et al., 2010) were shown to convert software and computing into metered services, with consequences for entrepreneurship, lock-in, and infrastructural concentration. Finally, the creator economy was analysed as platformized cultural production (Cunningham & Craig, 2019; Nieborg & Poell, 2018), with Duffy's (2017) aspirational labour supplying its social economics. The chapter's method is portable: for any digital business, ask who pays, who is subsidized, what is measured, and where the risk sits.

Discussion questions. (1) Take one service you use in both a free and a paid variant and analyse it three times: as a two-sided advertising business, as versioning, and as an acquisition funnel. Which analysis best predicts the firm's design choices? (2) Goldfarb and Tucker (2011) found that privacy regulation reduced advertising effectiveness. Does it follow that privacy and the free Internet are in necessary conflict? Specify the assumptions under which the inference holds and fails. (3) Using Hagiu and Wright (2015), advise a platform deciding between marketplace and reseller modes for a new category - and state what your advice implies for its third-party sellers. (4) 'The sharing economy shares nothing; it rents everything.' Assess, distinguishing asset-sharing, labour on demand, and professionalized commercial supply. (5) The cloud rebuilt the mainframe: evaluate this chapter's claim by comparing the economics of time-sharing in the first wave with hyperscale cloud in the fourth. (6) Design a creator's income portfolio across advertising shares, sponsorship, and direct monetization, and analyse each stream's exposure to platform policy risk. What would 'diversification' even mean for a platform-dependent entrepreneur?

5.5a. Numerical and Data Exercises

Exercise 5.A (SaaS unit economics). A subscription firm has ARPU EUR 50/month, gross margin 80%, monthly churn 4%, and CAC EUR 600. (a) Expected customer lifetime (months). (b) Lifetime value, LTV = ARPU × gross margin × lifetime. (c) Compute the LTV:CAC ratio and the CAC payback period; is the business healthy by the common 3:1 and "under 12 months" rules of thumb? (Compare the Worked Example in Section 5.4, which uses the same figures with 2% churn.)

Exercise 5.B (freemium). 500,000 free users; 3% convert to paid at EUR 10/month; each free user costs EUR 0.05/month to serve. (a) Monthly revenue from conversions. (b) Monthly cost of the free base. (c) Net monthly contribution, and the break-even conversion rate.

Indicative answers. 5.A (a) $1 \div 0.04 = 25$ months. (b) LTV = $50 \times 0.8 \times 25$ = EUR 1,000. (c) LTV:CAC = $1{,}000 \div 600 \approx 1.67$ (below 3:1); payback = $600 \div (50 \times 0.8) = 15$ months (above 12). The unit economics are weak - churn and/or CAC are too high. 5.B (a) 15,000 paid × EUR 10 = EUR 150,000. (b) 500,000 × EUR 0.05 = EUR 25,000. (c) Net = EUR 125,000; break-even conversion solves $500{,}000 \cdot c \cdot 10 = 500{,}000 \cdot 0.05 \rightarrow c = 0.5\%$.

5.6. Case Studies and Debates

The materials below dissect one composite business model in full, follow the subscription turn through the industry where its stakes are civic, examine a regional marketplace that held its ground against a global entrant, and stage the decade's signature dispute over platform tolls.

Case 5.1. The flywheel: anatomy of a composite business model

The large e-commerce platform that grew from an online bookstore into one of the world's most valuable firms is the indispensable specimen for this chapter, because it operates nearly every model in the taxonomy simultaneously and couples them deliberately. Its founding narrative is the *flywheel*: lower prices draw customers; customer traffic draws third-party sellers (the marketplace of Section 5.2); seller competition and scale lower costs and prices further; and the cycle turns. Onto this loop the firm grafted, layer by layer, the chapter's remaining models: a reseller business alongside the marketplace, with the mode boundary managed strategically (Hagiu & Wright, 2015); a subscription programme bundling delivery with media, engineered to raise purchase frequency and switching costs (Chapters 2 and 3). To these it added a logistics service selling fulfilment to its own sellers; an advertising business auctioning search placement - the aggregator model of Section 5.2 growing inside the marketplace, now a profit centre of the first rank; and the cloud division (Section 5.3), born of internal infrastructure, whose margins long subsidized the retail flywheel's thin ones.

The composite structure is analytically consequential, not merely impressive. Each layer feeds the others data and demand - the marketplace reveals what sells, the reseller and private-

label operations can act on the revelation (Zhu & Liu, 2018), the subscription locks in the demand side while logistics locks in supply, and advertising taxes sellers for visibility in rankings the platform controls: the conflicts of interest of Section 3.4 assembled into an architecture. It is precisely this coupling that Khan (2017) placed at the centre of the antitrust reform debate of Chapter 10 - harms legible as dependence, self-preferencing, and structural conflict rather than as consumer prices, which the model keeps conspicuously low. Seminar questions: (1) Map every revenue stream of the firm onto this chapter's taxonomy, and identify for each who pays, who is subsidized, what is measured, and who bears risk. (2) Which single layer, if structurally separated, would do most to resolve the conflicts of interest - and what efficiency would separation destroy? (3) The flywheel was drawn as a virtuous circle for consumers; redraw it from the standpoint of a third-party seller, using Chapters 3 and 4.

Case 5.2. The subscription turn in journalism: rescue and stratification

Chapter 9 will treat journalism's crisis as a media question; here it serves as the business-model experiment. When the advertising bundle collapsed (Section 5.1's history run in reverse), the strategic response of quality publishers was the subscription pivot: metered paywalls, digital-first products, and the conversion of occasional readers into recurring subscribers. The flagship successes are real and instructive - a handful of global titles built digital subscriber bases in the millions, diversified into games, audio, and consumer verticals (the bundle reborn in miniature), and now earn more from readers than from advertisers, realigning incentives exactly as Section 5.1's theory predicts: product decisions optimize perceived ongoing value - retention - rather than click-through.

The system-level result, however, is the stratification Section 5.1 flagged. Subscription economics reward publications with large addressable audiences of paying propensity - national and international titles serving educated, affluent readers - while local and regional journalism, whose civic function is least substitutable, commands neither the audience scale nor the willingness to pay, and continues to contract (Pickard, 2020). The paywall, moreover, prices exactly the verified information whose free circulation the misinformation debates of Chapter 9 presuppose: the reliable report costs money while the fabrication is free, a perverse structure from the standpoint of the information commons. Nielsen and Ganter's (2018) publisher-platform dependence completes the picture: even paywalled titles need platform distribution for discovery, so the subscription turn reduces but does not remove the asymmetry. Seminar questions: (1) Using lifetime value and churn, explain why subscription news favours scale and habit; which editorial consequences follow? (2) Design a business model for local news using

any combination of this chapter's instruments, and state honestly which market failure your design leaves standing for Chapter 10's policy instruments. (3) Is the 'reliable news costs, misinformation is free' structure a market failure in the Chapter 2 sense? Specify the good, the externality, and the remedy.

Case 5.3. Allegro: the regional marketplace that did not tip

Case 5.1 examined the flywheel of a global e-commerce platform. Allegro, Poland's largest online marketplace, shows how the same mechanisms operate when the incumbent is regional and the entrant is global. Founded in 1999 as an auction site, Allegro evolved into a fixed-price marketplace for third-party sellers, added its own retail operation, and introduced Allegro Smart!, a subscription bundling free delivery and returns - the purchase-frequency and switching-cost logic that Case 5.1 attributed to its global counterpart. Its 2020 listing on the Warsaw Stock Exchange was the largest initial public offering in the exchange's history. When the world's largest e-commerce firm opened a dedicated Polish storefront in 2021, many observers expected the market to tip.

It did not tip, and the reasons are those of Section 3.3. Network effects in a marketplace are partly local: Polish buyers value Polish sellers, Polish-language listings, local payment methods such as BLIK (Case 3.3) and dense parcel-locker delivery. The incumbent's installed base of sellers and its logistics arrangements were therefore not easily replicated by an entrant whose global scale did not transfer. Multi-homing was cheap for buyers but costly for sellers who had built reputation histories on Allegro, and those non-portable ratings (Chapter 8) acted as a switching cost.

The case has a second half that brings it back to Section 3.4 and Chapter 10. In December 2022 the Polish competition authority, UOKiK, fined Allegro more than PLN 206 million for favouring its own store. According to the authority, the platform had used information unavailable to independent sellers and reserved promotional tools for its own offers from 2015 onward - a textbook case of the self-preferencing that the Digital Markets Act later prohibited for designated gatekeepers (Prawo.pl, 2022). Allegro contested the decision. A regional champion, in other words, can reproduce the conflicts of interest of the global platforms to which it is presented as an alternative. Seminar questions: (1) Using the tipping conditions of Section 3.3, explain why the global entrant's scale did not translate into a winning position in Poland; which condition would you monitor for signs of change? (2) Is self-preferencing by a national marketplace less harmful than self-preferencing by a global gatekeeper? Consider market definition, the DMA's designation thresholds and national

competition law. (3) Redraw the flywheel of Case 5.1 for Allegro and mark the points at which local institutions - payments, logistics, language - enter the loop.

Debate 5.1. The thirty per cent question: app store commissions

Motion: dominant app stores' commissions and anti-steering rules constitute an abuse of gatekeeper position and should be capped or dismantled by regulation. The dispute is concrete: mobile ecosystems require most in-app digital purchases to flow through the platform's payment system at commissions historically around thirty per cent, and long prohibited developers from steering users to cheaper channels - terms contested in litigation across jurisdictions, most prominently *Epic Games v. Apple* (2023) - where a 2025 contempt order barred the platform from charging any commission on purchases made through external links (Epic Games, Inc. v. Apple Inc., 2025) - and addressed directly by the Digital Markets Act's anti-steering and sideloading provisions (Chapter 10). *For the motion:* the commission is a toll on a bottleneck, not a price formed by competition - developers cannot reach the installed base except through the store (the complementor's predicament of Section 3.4 in its purest form); anti-steering rules exist precisely to disable the price competition that would discipline the toll; the rate bears no visible relation to the cost of payment processing or review; and the platform exempts its own services, competing with complementors on terms it sets. *Against the motion:* the commission prices a bundle - distribution to a curated, secure, global market; payment, hosting, tools, and discovery - that developers previously assembled at far greater cost, and the counterfactual price of that bundle, not zero, is the benchmark. The store's security and quality controls are the trust machinery of Chapters 2 and 8, which sideloading mandates erode for all users to benefit a few large developers well able to run their own storefronts. Most developers pay reduced small-business rates, and the two-sided price structure (Section 3.2) may efficiently tax the money side - successful paid apps - to subsidize the free ecosystem users demonstrably value. Participants should be required to propose the *remedy* they would order - rate cap, mandatory alternative payments, sideloading, structural separation of store from platform - and defend it against the failure modes identified in Debate 3.1.

5.7. In Depth: Pricing the Free - The Machinery of Digital Advertising

Section 5.1 described the advertising-funded model from the business side; this advanced section opens the machinery that prices attention itself, because the auction, the attribution system, and the fraud economy beneath the free Internet are among the digital economy's most consequential and least examined institutions. The reader who works through this section will understand what is actually sold, by what mechanism, at whose risk - and why the machinery's measurement claims deserve the scepticism that Sections 6.3 and 9.1 apply to them.

The auction: pricing every impression separately

The pivotal institutional invention of the advertising Internet was the keyword auction. Where broadcast advertising sold audiences in bulk - a demographic estimate, a rate card, a season - search advertising priced each query impression separately, in a continuous auction among advertisers bidding for position. Edelman et al. (2007) supplied the canonical analysis of the mechanism that became the industry standard, the *generalized second-price* (GSP) auction: bidders are ranked (by bid, weighted in practice by predicted relevance), positions are awarded in rank order, and each pays approximately the bid of the bidder below - a design that keeps bids honest enough, revenues high, and the auctioneer's incentives interestingly mixed, since the relevance weighting means the platform's quality predictions co-determine prices. The economic significance is threefold. Granularity: pricing the individual impression converts the attention economy of Chapter 9 from metaphor to market microstructure - attention literally clears at auction, in milliseconds, billions of times daily, with the real-time bidding systems of display advertising extending the mechanism from search queries to essentially every page view, each accompanied by the tracking dossier of Chapter 2. Self-revelation: the auction extracts advertisers' private valuations continuously, so the platform learns the commercial value of every keyword, audience, and moment - data capital (Chapter 6) accruing from the pricing mechanism itself. And market power: the auctioneer sets the reserve prices, the quality weights, and increasingly - as both bidding automation and the exchanges consolidated into the largest platforms - operates on both sides of the market it clears, the conflict-of-interest structure of Sections 3.4 and 5.2 reproduced in market microstructure, and the subject of the advertising-technology antitrust actions contemporary with this book.

Attribution: the epistemology of the click

The advertising model's historic promise - finally, the measurable half of the marketing budget - rests on attribution: the assignment of observed outcomes (clicks, purchases, installations) to the advertisements that allegedly caused them. The epistemological difficulty is the one every empirical social scientist will recognize: targeting creates selection. The systems of Section 6.2 are optimized to show advertisements to people predicted to convert - which means the advertisement is systematically shown to people who would have converted anyway, and naive attribution (last click, view-through) credits the targeting for the conversion. The experimental literature that measured advertising's *incremental* effect - by randomized holdouts, the only clean identification - repeatedly found lifts far below what attribution dashboards reported, with the gap largest precisely for the most targeted campaigns: brand-keyword search advertisements harvesting users already en route, retargeting crediting itself for returns that loyalty produced (Blake et al., 2015; Lewis & Rao, 2015). The finding does not make advertising worthless; it makes the *measurement claims* of the extraction complex - the claims on which, as Section 6.3 notes, both the industry's pricing and its critics' alarm partly rest - systematically inflated, and it explains an otherwise puzzling fact: advertisers with sophisticated experimentation capacity periodically discover, and publicize, that large tranches of their digital spending were incremental to nothing. For the student, the lesson is a template: wherever the digital economy sells prediction-based performance, ask what the counterfactual is and who is allowed to measure it.

Fraud and the ecology of fake attention

A market that pays for attention summons counterfeit attention, and the fraud economy is the model's permanent shadow. Its inventory is instructive because each species exploits a specific link in the chain this section has assembled: bot traffic manufactures the impressions the auction prices; click farms manufacture the engagements attribution counts; domain spoofing sells the counterfeit of premium context; installation fraud harvests app-marketing budgets by simulating the conversions that bidding algorithms optimize toward - so that the fraud, once inside the measurement loop, is *amplified* by the very automation that was supposed to guarantee performance. Industry estimates of fraud's share of digital advertising spending vary by an order of magnitude - itself a datum about the measurement environment - and the countermeasure economy (verification vendors, private marketplaces, walled gardens whose internal metrics cannot be externally audited) reproduces the trust problematic of Chapters 2 and 8 at the heart of the model: the walled garden solves fraud by asking the advertiser to trust the seller's own scorekeeping, exchanging one verification problem for another. The synthesis

returns the section to the book's method - who pays, who is subsidized, what is measured, where does risk sit: in digital advertising, the advertiser pays, the user is subsidized, the platform measures, and the risk of counterfeit and of inflated counterfactuals sits, structurally, with the party least able to audit. Discussion questions: (1) In a GSP auction, explain how the quality weight changes both allocation and price, and identify the incentive it gives the auctioneer. (2) Design the randomized experiment that measures a retargeting campaign's incremental effect, and list the interests that would resist running it. (3) Why do walled gardens reduce fraud and reduce auditability simultaneously? Connect to the verification costs of Chapter 2. (4) If attribution systematically overstates advertising's effect, what follows for the surveillance-capitalism debate of Section 6.3 - and what, carefully stated, does not follow?

5.8. In Depth: Play Money - The Economics of Games and Virtual Goods

The games industry outsells recorded music and cinema combined, pioneered most of the monetization techniques this chapter catalogued, and runs, inside its products, functioning economies complete with currencies, inflation, and labour markets; this section takes it seriously as economics, following the field Castronova (2005) founded and Lehdonvirta and Castronova (2014) systematized.

Virtual scarcity: the deliberately un-digital good

Chapter 2 derived the digital economy from non-rivalry and zero marginal cost; the virtual good is the instructive exception, because its scarcity is not overcome but *manufactured*. The sword, the skin, the plot of virtual land could be duplicated at zero cost; the operator declines to, because the good's value is positional and social - rarity is the product - and so writes scarcity into code: drop rates, mint limits, untradeable bindings. Castronova's (2005) founding observation was that these designed scarcities generate real economics - measurable exchange rates against dollars, gross products per capita for game worlds, and cross-border labour (the gold-farming workshops that were ghost work, Section 7.3, a decade early). Lehdonvirta and Castronova (2014) drew the discipline's lesson: virtual economies are macroeconomic *design problems*, in which the operator sets money supply (faucets and sinks), price levels, and property rules, and faces, in miniature and with perfect data, the policy trade-offs of a central bank crossed with a land registry. The theoretical payoff runs both directions: the game economy is Chapter 2's theory with the parameters exposed - proof, if more were needed, that scarcity in the digital economy is always an institutional choice - and the game designers' solved problems (sinks against inflation, binding against inequality spirals) prefigure governance questions that 'serious' digital assets rediscovered expensively.

Free-to-play and the frontier of monetization ethics

The industry's business-model history compresses this chapter into two decades: boxed product (the packaged software of Section 5.3's prehistory), subscription worlds, and then the free-to-play settlement - freemium (Section 5.1) perfected, with the game free, the whales financing it, and monetization moved inside the play loop as cosmetic goods, battle passes (subscription re-entering as seasonal commitment), and randomized loot boxes. The model's economics are Section 2.2's versioning plus a demand-side engine the other industries lack: the goods are positional within a community the game itself convenes, so the platform manufactures both the scarcity and the audience before whom scarcity signals. Its controversies

are correspondingly the chapter's sharpest: loot boxes price a probability, not a good - some jurisdictions (most firmly Belgium) treat paid loot boxes as gambling, others mandate probability disclosure (a certification regime, Section 2.8) - and the engagement engineering of Section 9.1 operates here on players including minors, making games the live laboratory for the design-ethics regulation that Chapter 10's consumer-protection frontier is only beginning to codify. The section's closing bridge is to the property question: player-owned tradeable assets - whether on the operator's ledger or on public blockchains - re-pose Case 6.2's question one layer down: what does ownership mean when the asset's existence, scarcity, and usability are revocable rules in someone else's system? Discussion questions: (1) Explain why an operator profits from *not* duplicating a costlessly duplicable good; connect positional goods to Section 2.2's versioning. (2) Design the monetary policy of a game economy: name two faucets, two sinks, and the inflation indicator you would target. (3) Should loot boxes be regulated as gambling, as a disclosure problem, or not at all? Argue from mechanism, not vibes. (4) Assess 'true digital ownership' claims for tradeable game assets using Chapter 2's excludability analysis and Section 3.7's platform-as-regulator.

5.9. In Depth: Money as a Platform - Payments, Open Banking and Digital Currencies

Payments are where the platform economics of Chapter 3 meet the oldest network good of all - money. This advanced section treats the payment system as a layered platform, examines how data-access regulation has tried to open it, and asks what digital currencies issued by central banks or private firms would change.

Payments as a two-sided network

Every payment instrument is a two-sided market in the sense of Rochet and Tirole (2003), whose model was built on card schemes. Consumers value an instrument that many merchants accept, merchants value one that many consumers carry, and the scheme sets a price structure - interchange fees, rewards, merchant charges - across the two sides. Digitalization multiplied the layers of this platform. Wallets and super-apps sit on top of cards and bank accounts and capture the customer interface. Account-to-account schemes such as India's UPI, Brazil's Pix and Poland's BLIK (Case 3.3) bypass cards altogether. Mobile money in the tradition of M-Pesa (Case 4.1) creates stored value outside banks. The strategic question is always who owns the interface at the moment of payment, because the interface owner observes the transaction data and can envelop adjacent services - credit, insurance, loyalty - in the manner of Section 3.4.

Open banking: mandated data portability in finance

The European Union's second Payment Services Directive (PSD2) turned the switching-cost analysis of Section 3.8 into financial regulation: banks must allow licensed third parties, with the customer's consent, to access account information and to initiate payments (European Parliament & Council of the European Union, 2015). Open banking is thus data portability with an interface attached, and its effects follow the mechanism rather than the rhetoric. Babina et al. (2025), combining evidence on open-banking policies in 49 countries with UK microdata, find that such mandates spur fintech entry and widen access to financial advice and credit, including new lending relationships for small firms. When shared data are used mainly to screen credit risk rather than to improve products, however, higher-risk and privacy-conscious consumers can face worse terms. The finding restates the economics of privacy of Section 6.8 in a regulated industry: whether data flows help or harm depends on what they are used for. The EU's later rules on data sharing, above all the Data Act (Section 10.2), extend the same logic beyond payments.

Digital currencies: private stablecoins and public money

Two further developments challenge the bank-centred architecture. Stablecoins - privately issued tokens pegged to a sovereign currency and backed by reserve assets - promise cheap, programmable transfers, but they reproduce an old problem of private money: their value rests on the quality and liquidity of the reserves. This is why the EU's Markets in Crypto-Assets Regulation (MiCA) subjects issuers to authorization, reserve and redemption requirements (European Parliament & Council of the European Union, 2023a). Central bank digital currencies (CBDCs) approach the same space from the public side. Auer et al. (2022) summarize the motives - preserving access to public money as cash declines, fostering competition in payments, and reducing dependence on foreign or private infrastructures - and the risks, above all that deposits might migrate from banks to central-bank money in a crisis. The European Commission's 2023 proposal for a digital euro combines these motives with an explicitly strategic one: a public, pan-European means of payment that would reduce reliance on non-European card schemes and wallets (European Commission, 2023). Discussion questions: (1) Identify, for one payment method you use, the interface owner, the money side and the subsidy side. (2) Why might open banking increase competition yet harm some consumers? Use Babina et al. (2025) and Section 6.8. (3) Is a digital euro a competition policy, a monetary policy or an industrial policy? Construct the case for each reading.

Chapter 6. Data, Algorithms, and Artificial Intelligence as Economic Resources

Learning objectives. After studying this chapter, the reader should be able to: (1) explain in what sense data functions as capital in the digital economy, and evaluate the competing metaphors - oil, exhaust, labour, colony - through which its extraction is described; (2) present the economics of prediction and explain why machine learning diffuses as a general-purpose input; (3) analyse algorithmic decision-making as an institution, including the problems of opacity, bias, and accountability; (4) reconstruct Zuboff's theory of surveillance capitalism together with its principal criticisms; and (5) frame the economics of generative artificial intelligence within the automation debates and the productivity-paradox tradition of Chapter 1.

KEY CONCEPTS

data capital; prediction; algorithmic decision-making; foundation models; inference; compute; model API; open weights; AI-native firm; agents; surveillance; privacy; cyber resilience

Every chapter so far has brushed against data - as the product of cheap tracking (Chapter 2), the fuel of platform feedback loops (Chapter 3), the telemetry of lean experimentation (Chapter 4), and the input of targeted advertising (Chapter 5). This chapter makes the resource itself the object of analysis, together with the algorithmic systems that convert it into predictions, decisions, and, most recently, generated content. The stakes of getting the analysis right are high on all three of the book's registers: economically, data and AI are candidate general purpose technologies in the sense of Chapter 1; sociologically, they institute new forms of classification and control; and for media scholarship, they constitute the selection machinery through which public communication now flows.

6.1. Data as Capital: Extraction, Aggregation, and Valuation

The popular metaphor calls data 'the new oil'; the scholarly literature has spent a decade explaining why the metaphor misleads. Oil is rival, depletable, and valuable in the barrel; data is non-rival (Chapter 2), self-renewing wherever activity is instrumented, and nearly worthless in the datum - its value arises only through aggregation, linkage, and analysis, and it is highly context-specific: the clickstream that is gold to an advertising business may be dross to an insurer. Mayer-Schönberger and Cukier (2013) popularized the underlying transformation as *datafication*: the rendering of ever more social action - movement, speech, attention, sociality itself - into quantified, machine-readable form. Sadowski (2019) supplies the sharper political-economic frame: data is best understood not as a commodity sold on markets - most of it never is - but as *capital*: an asset accumulated for its capacity to generate value, whose pursuit reorganizes firm behaviour around extraction. The capital frame explains observable conduct that the commodity frame cannot: firms collect data far beyond present use because accumulation is the point; products are redesigned as collection instruments; and acquisitions are priced on data assets that appear on no balance sheet. Valuation is correspondingly unsettled - accounting standards capitalize databases poorly, national accounts (Chapter 1's measurement problem) miss data assets almost entirely, and the gap between the market capitalization of data-rich firms and their book value is, in part, the shadow price of an asset the statistical system cannot see.

Who, then, owns the asset? The legal answer is deliberately oblique - data is governed by a lattice of privacy, contract, and database rights rather than by property - and the normative debate has produced a spectrum of metaphors, each encoding a programme. ‘Data as labour’ proposes that users are unpaid producers who should be compensated; ‘data as exhaust’ naturalizes collection as the harmless byproduct of service use; Couldry and Mejias (2019) propose *data colonialism*: the appropriation of human life through data, they argue, repeats the colonial move of declaring a domain - once land, now lived experience - terra nullius available for capture, with the platform firms of Chapters 3 and 5 as the appropriating powers. One need not adopt the strongest metaphor to accept its analytical core, which this book takes as established: the data economy rests on a founding asymmetry - individuals generate the resource, firms capture and capitalize it - and the terms of that capture were set by default and by design rather than by bargain. The regulatory responses, from consent regimes to portability mandates, are examined in Chapter 10.

6.2. Algorithmic Decision-Making and Prediction Machines

What does the accumulated capital do? The parsimonious answer of Agrawal et al. (2018) has become the standard economic frame: machine learning is, economically, a drop in the cost of *prediction* - the filling-in of missing information from available information. When prediction becomes cheap, three consequences follow with textbook logic. Activities that are explicitly predictive - demand forecasting, credit scoring, fraud detection, targeting - expand and improve. Activities not previously organized as prediction are reframed to use it: translation, diagnosis, driving, hiring, content ranking. And the complements of prediction appreciate: data, on which predictions train, and human *judgment* - the specification of objectives and the weighing of what the machine cannot value - which the framework treats as the durable human contribution. Varian (2014) makes the methodological companion point: the tools of machine learning entered economics and business as instruments for exactly the large, granular datasets that datafication supplies. The prediction frame also explains diffusion: like the general purpose technologies of Chapter 1, cheap prediction is pervasive, improving, and complementary to downstream innovation - and, true to the tradition of David (1990), its productivity effects arrive only with complementary organizational investment, a point to which Section 6.4 returns.

Economics explains why algorithmic decisions spread; sociology and law ask what kind of institution they are. Three problems structure the critical literature. The first is *opacity*: Pasquale (2015) describes the ‘black box society’, in which decisions of consequence - credit,

insurance, employment screening, visibility in the feeds of Chapter 9 - are taken by systems whose logic is shielded by trade secrecy and, in complex models, by genuine technical inscrutability, inverting the classical relation in which the governed can inspect the rules that govern them. The second is *bias*: Barocas and Selbst (2016) demonstrate that discrimination requires no discriminatory intent - models trained on historical data learn historical hierarchies; target variables and proxies encode disadvantage (the postcode that stands in for race, the employment gap that stands in for motherhood); and the resulting disparate impact fits poorly the intent-centred architecture of anti-discrimination law. O'Neil (2016) adds the dynamic pathology: opaque, scaled, and self-confirming models - her 'weapons of math destruction' - create feedback loops in which the predicted risk produces the treatment that produces the outcome that vindicates the prediction. The third is *accountability*: when a consequential decision is distributed across data collection, model design, deployment choices, and automated execution, responsibility diffuses - the problem that European regulation attacks with transparency and human-oversight requirements discussed in Chapter 10. The synthesis matters for this book's argument: algorithmic decision-making is not merely a technology adopted by institutions; it *is* an institution - a way of classifying persons and allocating outcomes - and it should be analysed with the tools sociology reserves for institutions, a thread Chapter 8 picks up under the heading of quantification.

6.3. Surveillance Capitalism and Its Critics

The most influential synthesis of Sections 6.1 and 6.2 is Zuboff's (2019) theory of *surveillance capitalism*. Its claims can be stated in the book's accumulated vocabulary. Firms discovered, in the advertising economics of Chapter 5, that the data exhaust of free services could be refined into *behavioural surplus* - data beyond what service improvement requires - and sold as prediction products in what Zuboff calls behavioural futures markets, of which the real-time advertising auction is the prototype. Competition on prediction quality drives an *extraction imperative* - ever more data, from ever more domains of life, hence the instrumented home, car, and body - and, at the limit, an intervention imperative: the most valuable prediction is the one the predictor helps come true, so systems graduate from forecasting behaviour to shaping it, through the engagement engineering that Chapter 9 examines. The theory's force lies in its unification: what look like separate controversies - privacy erosion, addictive design, platform power - appear as expressions of a single accumulation logic, as characteristic of its epoch as Fordist mass production was of its.

Precisely because the theory has been so influential, its criticisms deserve equal statement. Political economists in the tradition of Srnicek (2017) object that surveillance is a strategy of *some* platform businesses, not the essence of digital capitalism - the cloud, logistics, and enterprise-software giants of Chapter 5 accumulate without behavioural futures markets - and that Zuboff's implied golden age of pre-surveillance capitalism flatters what was, on any reading, already an audience-selling media economy (Smythe, 1977). Empiricists note that the advertising machinery's persuasive power may be substantially overstated - targeting lifts are real but modest in careful studies (Blake et al., 2015; Lewis & Rao, 2015), and an industry that sells influence has every incentive to exaggerate it - so that the theory risks taking the extraction complex's marketing at its word. And scholars of data colonialism (Couldry & Mejias, 2019) argue the frame is too narrow rather than too broad: centring manipulation of consumers in rich markets understates the global and infrastructural character of data appropriation. The methodological lesson for the student is the book's recurring one: hold the mechanism (which is well documented - the extraction architecture exists) apart from the strongest claims about its efficacy (which remain contested), and notice that regulation must act, as Chapter 10 shows, under precisely this uncertainty.

6.4. Generative AI and the New Automation Debate

The fifth wave of Chapter 1 entered public consciousness when machine learning systems began to *generate* - text, images, code, audio - at a quality sufficient to substitute, in some tasks, for skilled human output. Economically, generative AI extends the prediction frame of Section 6.2 to its logical completion: the predicted missing information is now the artefact itself. Its labour-market signature differs instructively from previous automation. Where the technologies analysed in Chapter 7 substituted most readily for routine and manual tasks, exposure analyses of large language models find the highest task overlap in cognitive, white-collar, and creative work - writing, coding, analysis, design - inverting the historical gradient (Eloundou et al., 2024). Whether exposure becomes displacement, augmentation, or new-task creation is the substance of the automation debate proper, for which Acemoglu and Restrepo (2019) supply the essential accounting: automation's displacement effect competes with productivity effects, capital deepening, and - decisively in the historical record - the *reinstatement* effect of newly created tasks in which labour has comparative advantage; the net outcome is not a fact about the technology but a resultant of these forces and of the institutions that steer them.

COMMON MISCONCEPTION

Exposure to generative AI is not equivalent to job displacement. Exposure describes the share of tasks whose execution could be altered by a technology; employment effects depend on adoption, complements, organizational redesign, demand responses, regulation, and the creation of new tasks.

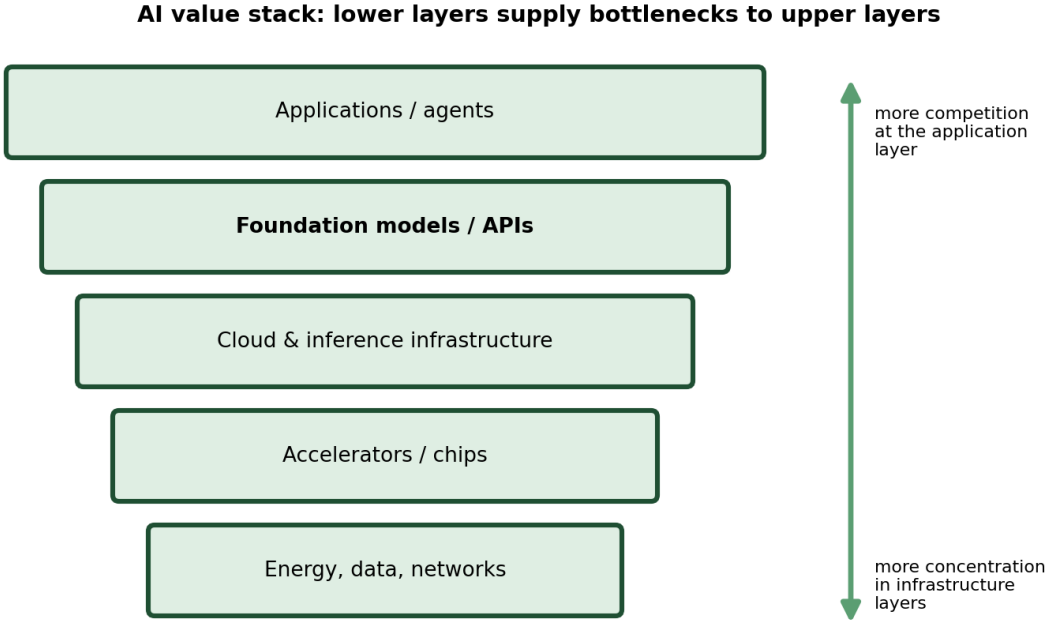

Figure 6.1. The AI value stack: applications depend on increasingly concentrated infrastructure layers.

The productivity question completes the circle to Chapter 1. Brynjolfsson, Rock, and Syverson (2019) diagnose a *modern productivity paradox*: expectations for AI are enormous while measured productivity growth remains sluggish, and the most plausible reconciliation is the implementation-lag account inherited from David (1990) - general purpose technologies deliver only after complementary investment in processes, skills, and organization, which takes decades and is itself costly and unmeasured. The honest state of knowledge at this book's writing is therefore symmetrical: the historical priors say neither mass unemployment nor effortless abundance, but a protracted, institutionally mediated reallocation whose distributional consequences - which tasks, whose wages, which regions - are the true policy variables. Those consequences are the business of the next chapter; the governance of the systems themselves returns in Chapter 10.

The first wave of micro-level evidence sharpens this picture without resolving it. Experimental and quasi-experimental studies find sizeable task-level gains. Professionals given access to a chatbot completed writing tasks in markedly less time and at higher assessed quality (Noy & Zhang, 2023), and customer-support agents using a generative assistant resolved about 14 per cent more issues per hour, with the largest gains for the least experienced workers

(Brynjolfsson et al., 2025). Two findings keep the macroeconomic verdict open. Acemoglu (2025), aggregating task-level estimates over the share of tasks plausibly affected, predicts total factor productivity gains of well under one per cent over a decade. Humlum and Vestergaard (2025), using Danish administrative data, find no detectable effect on earnings or recorded hours in the first two years after chatbots spread, even as work inside firms was being reorganized. The pattern is the implementation lag of David (1990) observed in real time: tasks change first, measured output later, if at all.

Applied Extension: The Economics of Generative AI and AI-Native Firms

Generative AI changes the entrepreneurial problem twice. First, it lowers the cost of producing cognitive artefacts - text, code, images, prototypes, analysis - so that very small teams can test and serve markets that previously required specialized labour. Second, it introduces a new dependency stack: applications may be cheap to build while relying on foundation-model providers, hyperscale clouds, accelerators, energy and data infrastructures whose economics are substantially more concentrated. Entry barriers can therefore fall at the application layer while rising at the infrastructure layer.

An AI-native venture should be analysed along five margins: model access (proprietary API, open-weight model, or in-house model); inference cost per useful task; data advantage and rights to use that data; distribution and workflow integration; and the degree to which the product creates proprietary feedback rather than merely wrapping a commodity model. The strategic question is not simply whether the underlying model is powerful, but whether the venture owns a scarce complement - trusted distribution, workflow position, proprietary data, domain expertise, evaluation infrastructure, or regulated access - that remains valuable as model capability commoditizes.

Unit economics must also include inference explicitly. A product with near-zero software replication costs may still face material variable costs when each user request consumes model inference, retrieval, verification, human review or third-party tools. The relevant margin is therefore contribution after inference rather than gross software margin in the classical SaaS sense. For high-stakes applications, the cost of evaluation, monitoring, audit and human escalation belongs in the production function rather than being treated as an external compliance overhead.

A further shift is under way from generative systems that answer to agentic systems that act: software agents that plan multi-step tasks, call tools and application programming interfaces, and transact on a user’s behalf - booking, purchasing, negotiating, writing and executing code. Economically, agents push the collapse of search and transaction costs one step

beyond recommendation (Section 2.7), because the agent rather than the consumer compares the offers. This relocates the aggregator position of Section 5.2: whoever controls the default agent can steer demand across merchants and complementors, and the self-preferencing concerns of Section 3.4 and Chapter 10 return in a new form. Agents also raise liability questions that existing law answers poorly - who bears the loss when an agent makes a mistaken purchase or an unlawful commitment? Chan et al. (2023) argue that the harms of such systems grow with their autonomy, goal-directedness and capacity for long-horizon planning, and that governance should anticipate these properties rather than react after the fact. For the entrepreneur, the lesson of this section applies with extra force: an agent built on another firm's model and distributed through another firm's channel owns little unless it controls a scarce complement.

WORKED EXAMPLE

Scenario. An AI service performs 2 million inference calls per month. Average model cost is EUR 0.012 per call, while the product earns EUR 0.035 contribution revenue per call before fixed costs.

Calculation. Inference cost = EUR 24,000. Contribution revenue = EUR 70,000. Contribution after inference = EUR 46,000, or EUR 0.023 per call.

Sensitivity. A 50% rise in inference cost lowers contribution to EUR 34,000; a 20% fall in price lowers it to EUR 32,000. AI-native business models therefore depend on both model capability and the cost curve of serving each task.

THREE LENSES

Economics asks: Where are scarcity, scale economies, complements, and rents located across data, compute, models, and applications?

Sociology asks: How do classification systems redistribute autonomy, opportunity, surveillance, and accountability?

Media studies asks: How do generative systems change authorship, provenance, discoverability, and the circulation of synthetic content?

2026 UPDATE: AI ADOPTION, INDUSTRIALIZATION, AND INFRASTRUCTURE

Firm-level AI use has moved from experimentation toward broader organizational adoption. OECD data for countries with comparable statistics report AI use by 20.2% of firms in 2025, up from 8.7% in 2023, with a substantial size gap between large and small firms (OECD, 2026).

The 2026 AI Index documents continued capability gains, rapid private investment, expanding agentic systems, and increasing industrial concentration in model development and infrastructure (Stanford Institute for Human-Centered Artificial Intelligence, 2026).

The physical bottleneck is equally important: the IEA reports rapidly rising electricity demand from data centres and emphasizes grids, power equipment, advanced chips, and cooling as constraints on AI expansion (IEA, 2025, 2026).

Analytical implication: the economics of AI should be read as a joint problem of software, compute, energy, capital, organizational complements, and regulation - not as a model-performance story alone.

6.5. Chapter Summary and Discussion Questions

This chapter analysed the newest strata of the digital economy's resource base. Data was theorized as capital rather than commodity (Sadowski, 2019), datafication defined (Mayer-Schönberger & Cukier, 2013), and the ownership debate mapped across its governing metaphors, including data colonialism (Couldry & Mejias, 2019). Machine learning was framed as a fall in the cost of prediction with data and judgment as appreciating complements (Agrawal et al., 2018; Varian, 2014), and algorithmic decision-making was analysed institutionally

through opacity (Pasquale, 2015), bias without intent (Barocas & Selbst, 2016; O'Neil, 2016), and diffused accountability. Zuboff's (2019) surveillance capitalism was reconstructed as a unifying accumulation logic and confronted with its political-economic and empirical critics. Finally, generative AI was placed within the automation accounting of Acemoglu and Restrepo (2019), its inverted exposure gradient noted (Eloundou et al., 2024), and its productivity puzzle assigned to the implementation-lag tradition (Brynjolfsson, Rock, & Syverson, 2019).

Discussion questions. (1) 'Data is the new oil.' Enumerate the disanalogies and propose the metaphor you find most analytically productive, defending the programme it encodes. (2) Using the prediction framework, decompose a profession you know into prediction and judgment tasks; which components appreciate and which depreciate as prediction cheapens? (3) Barocas and Selbst (2016) argue that anti-discrimination law is mismatched to algorithmic disparate impact. Reconstruct the mismatch and evaluate one remedy: forbidding protected variables, auditing outcomes, or mandating explainability. (4) State surveillance capitalism's central mechanism in the vocabulary of Chapters 2, 3, and 5, then state the strongest empirical objection to it. Can both be right? (5) Why might generative AI's productivity effects be currently invisible? Compare the mismeasurement and implementation-lag explanations, and specify evidence that would separate them. (6) Acemoglu and Restrepo distinguish displacement from reinstatement. Identify one historical episode of each, and argue which force current institutions favour.

6.5a. Numerical and Data Exercises

Exercise 6.A (the economics of prediction). Following Agrawal et al. (2018), AI makes prediction cheap. A firm can use analysts (EUR 40 per prediction) or build an ML model (EUR 2,000,000 fixed, then EUR 0.001 per prediction). (a) Above what volume is the model cheaper? (b) Why does cheap prediction raise the value of complements such as judgment and data? (c) Give one distributional concern.

Indicative answers. 6.A (a) Set $40q = 2{,}000{,}000 + 0.001q \rightarrow q \approx 50{,}000$ predictions to cover the fixed cost; above that the model is cheaper. (b) When prediction is cheap, the scarce, valuable inputs are deciding what to predict and how to act on it (judgment) and the data that trains the model; as substitutes for prediction fall in price, these complements rise in value. (c) The fixed cost and data requirements favour large incumbents (concentration); prediction tasks may be automated, displacing some workers; and biased training data can entrench discrimination.

6.6. Case Studies and Debates

The materials below take the chapter's arguments to their hardest terrain: an algorithm that scores liberty, a collision between generative models and the property settlement of Chapter 2, and the proposal to pay people for their data.

Case 6.1. Scoring recidivism: the fairness impossibility

The controversy that made algorithmic bias a public question concerned risk-assessment instruments used in American criminal justice to score defendants' likelihood of reoffending, informing decisions from bail to sentencing. Investigative analysis of one widely deployed instrument (Angwin et al., 2016) reported a racial asymmetry in its errors: among defendants who did not reoffend, Black defendants were roughly twice as likely as white defendants to have been scored high-risk, while among those who did reoffend, white defendants were more likely to have been scored low-risk. The vendor's rejoinder invoked a different fairness criterion: the scores were *calibrated* - a given risk score corresponded to approximately the same reoffence rate regardless of race - so the instrument, on its own test, treated the groups alike.

The scholarly resolution of this dispute is the case's lasting contribution, and it is a theorem rather than a verdict. Chouldechova (2017) proved that when base rates of the predicted outcome differ between groups, an instrument cannot simultaneously satisfy calibration and equality of false-positive and false-negative rates: the fairness criteria invoked by the critics and the vendor are mathematically incompatible except in degenerate cases, so *some* disparity is unavoidable and the choice among disparities is normative, not technical. Kleinberg et al. (2018) draw the institutional moral this book endorses: the algorithm did not create the trade-off - human judges face it too, invisibly - but it made the trade-off explicit, auditable, and therefore governable; the scandal of the score is partly the transparency the score forces upon practices that were always discriminatory in the dark. The case thereby sharpens, rather than settles, Section 6.2: bias without intent is real (Barocas & Selbst, 2016); perfect fairness is undefined; and the operative questions become who chooses the error trade-off, under what procedure, with what right of contestation - questions of institutional design, which Chapter 10's due-process requirements begin to answer. Seminar questions: (1) State the two fairness criteria formally and explain, intuitively, why differing base rates force the conflict. (2) 'The algorithm is an audit of the society that trained it.' Defend and attack. (3) Design the appeal procedure a defendant scored high-risk should have; which elements does the black-box problem (Pasquale, 2015) obstruct, and how would you compel them?

Case 6.2. Generative models and the training-data wars

Generative AI has forced the property settlement of Chapter 2 into renegotiation, because the technology's supply chain begins with the appropriation, at scale, of the open web's expressive content. The models of Section 6.4 are trained on corpora containing books, journalism, images, and code under copyright; rights-holders across those industries have sued the model developers, contending that ingestion is reproduction requiring licence and that model outputs can substitute for - and occasionally regurgitate - the originals; developers respond that training extracts statistical patterns rather than expression, constituting transformative fair use, and that a licensing requirement covering the entire training distribution is practically equivalent to prohibition. The litigation, still ongoing across jurisdictions as this book is written, has begun to produce doctrine - early United States rulings in 2025 diverged, one court treating training on lawfully acquired books as fair use while rejecting that defence for pirated copies (Bartz v. Anthropic PBC, 2025), another rejecting fair use where the trained system served a directly competing product (Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc., 2025); the analytical structure, however, is already legible with Chapter 2's instruments and is the reason the case belongs here.

The dispute is Arrow's (1962) trade-off staged between two creative populations rather than between creators and users. Landes and Posner's (1989) balance - incentives to create against the costs imposed on future creators who build from existing work - now has a machine on its second side: the model is the limiting case of cumulative creation, built from everything, and the doctrine that emerges will set the terms on which the data capital of Section 6.1 may be assembled from other people's expression. The market, meanwhile, is not waiting for doctrine: licensing deals between model developers and large publishers are constructing, transaction by transaction, a market solution to Arrow's disclosure problem - while atomized creators, lacking the bargaining scale of Chapter 9's publishers, face the familiar asymmetry of Chapter 3's complementors. Seminar questions: (1) Map the litigants' positions onto the static-versus-dynamic efficiency trade-off of Section 2.1; what does each side's victory do to each incentive? (2) Compare three regimes - free training, collective licensing, opt-out with machine-readable reservation - by their transaction costs and their distributional consequences for small creators. (3) Does the music settlement of Case 2.1 offer a template here, or do the goods differ in a way that breaks the analogy?

Debate 6.1. Should we be paid for our data?

Motion: individuals should be compensated for the data they generate - data should be treated as labour, not as exhaust. The affirmative programme was given its economic form by Arrieta-Ibarra et al. (2018): the machine-learning economy's decisive input is human-generated data, currently expropriated at a price of zero under the 'exhaust' framing of Section 6.1; treating data as labour would channel income to its producers, create incentives for higher-quality contribution, and correct a labour market in which, as AI output grows, the humans whose data trained it are recorded as economically idle - with data unions supplying the countervailing bargaining power that atomized sellers of anything have always lacked. *Against the motion:* the marginal value of any individual's data is, for precisely the aggregation reasons of Section 6.1, near zero - the value lives in the corpus and the model, so honest data wages would be trivial and the proposal's redistributive promise illusory. Pricing data legitimates and entrenches the extraction architecture that privacy law (Chapter 10) seeks to constrain - one does not fix surveillance by unionizing it. Transaction costs of metering, attribution, and payment would consume the surplus, and the genuinely scarce contributions - expert annotation, professional content - are already priced in the ghost-work and licensing markets of Chapter 7 and Case 6.2. The debate should be structured around the disagreement's hinge: whether data's value is attributable to individuals at all, or is intrinsically collective - and each side must specify the institutional machinery (data unions, trusts, collective licensing, or prohibition) its answer implies.

6.7. In Depth: The Production Economics of Machine Learning

Sections 6.2 and 6.4 treated machine learning from the demand side - what cheap prediction does to the economy; this advanced section treats the supply side - how prediction is produced, at what cost structure, and with what consequences for industrial organization. The stakes are the book's recurring ones: whether the AI layer of the digital economy inherits, escapes, or intensifies the concentration dynamics of Chapter 3.

The nonrivalry theorem: why data economics is not oil economics, formally

Section 6.1 dismantled the data-as-oil metaphor informally; Jones and Tonetti (2020) supply the formal treatment that belongs in any advanced course. Modelling data as a nonrival input - usable simultaneously by any number of firms without depletion, exactly the property Chapter 2 established for information - they derive a result with sharp policy content: because nonrival goods create value in proportion to the breadth of their use, the *allocation of property rights over data* has first-order welfare consequences. When firms own the data their operations generate, they hoard: each firm restricts use to itself, out of competitive interest and (the model's subtle point) out of concern for creative-destruction - shared data breeds rivals - so the economy forgoes the gains of broad use. When *consumers* own their data and can sell it widely, the nonrivalry is harvested: data flows to many uses, and the model's welfare ranking places consumer ownership above firm ownership, with outright bans on data trade - the privacy-maximalist corner - costly for the same reason hoarding is. The theorem disciplines both sides of this chapter's debates: against the extraction status quo, it shows that the current allocation (de facto firm ownership) is not merely unfair but inefficient, leaving nonrival value unexploited; against simple prohibition, it shows that the remedy for misallocated data is reallocation - portability, sharing mandates, data trusts, the Chapter 10 toolkit - rather than enclosure in reverse. It also gives Debate 6.1 its missing analytical floor: the case for paying individuals is weakest as wage arithmetic and strongest as an ownership regime that changes what the data is allowed to do.

Scale, compute, and the capital structure of intelligence

The defining empirical regularity of the frontier-model era is the *scaling relationship*: within the current paradigm, model capability improves smoothly and predictably with the joint scaling of parameters, training data, and computation - a production function discovered by measurement rather than design, whose consequence is that frontier capability became, above all, a capital expenditure. Training runs are priced in the tens and hundreds of millions of dollars

of computation; the computation runs on specialized processors supplied by a strikingly concentrated hardware layer; the processors live in the hyperscale data centres of Section 5.3, whose energy and siting requirements have made electricity the newest constraint in the stack; and the training corpora are the contested appropriation of Case 6.2. The industrial-organization reading follows directly from Chapter 3's toolkit: enormous fixed costs and near-zero marginal costs of serving an additional query (Chapter 2's structure at new magnitude); vertical structure - chips, cloud, models, applications - in which the cloud incumbents occupy the choke layer and, through investment and partnership, integrate into the model layer above them (the envelopment logic of Section 3.4 running upward); and data feedback loops in which deployment generates the usage data that improves the next model. Whether this yields durable concentration is the live question Section 10.4 poses, and the countervailing evidence deserves equal statement: model quality has diffused rapidly, open-weight models trail the frontier by months rather than years, algorithmic efficiency gains repeatedly cut the compute cost of any fixed capability, and distillation lets small models inherit large ones' abilities. The moat may therefore lie less in the model, a depreciating asset, than in the familiar places: distribution, data, switching costs, and the enterprise relationships of Chapter 5.

The supply chain of judgment: humans in the loop, priced

The production function has a human column that the scaling narrative obscures and Chapter 7 examines in detail. Contemporary model production consumes, alongside computation, industrial quantities of human judgment: data annotation and cleaning (the ghost work of Gray and Suri, 2019, now the upstream of the fifth wave). It also consumes preference labelling, in which workers rank model outputs to steer training toward helpfulness and away from harm - the mechanism through which the values of Section 6.2's 'judgment' are literally purchased at piece rates. To these come specialist red-teaming and evaluation and, at the supply chain's hardest edge, the moderation labour of Roberts (2019) applied to training data and outputs, with its documented psychological costs. The economics are those of Case 7.2 - planetary labour markets, requester power, invisibility by design - with an added irony the reader should hold onto: the systems marketed as the automation of judgment are trained on judgment bought cheaply from the same global workforce the automation discourse renders obsolete. The distributional accounting of Section 6.4 therefore has a production-side entry: whatever generative AI does to the demand for skilled work downstream, its supply chain has already reorganized a stratum of it upstream - at prices set by the bargaining structure of Chapter 7, and largely outside the statistics of Section 1.10. Discussion questions: (1) State the Jones-

Tonetti argument in one paragraph, and derive from it a concrete design for data portability that a regulator could implement. (2) Using Chapters 2 and 3, argue both sides: 'frontier models are a natural monopoly' versus 'models are a depreciating commodity and the moat is distribution.' Which observable trends over the next three years would decide it? (3) Map the human labour inputs of a deployed model onto Chapter 7's categories, and identify where each sits in the classification struggle of Section 7.3. (4) Electricity has entered the production function of intelligence; trace the consequences through Section 1.10's measurement conventions and Section 10.3's geography.

6.8. In Depth: The Economics of Privacy

Sections 6.1-6.3 treat privacy as an object of extraction, and Chapter 10 treats it as an object of regulation; this section supplies the economic theory of the thing itself, following the field-defining survey of Acquisti et al. (2016), because the privacy debates of the digital economy are conducted, on all sides, in implicitly economic terms - trade-offs, valuations, harms - that repay being made explicit.

Two traditions and a synthesis

The economics of privacy inherited two opposed intuitions. The Chicago tradition treated privacy as concealment: personal information is relevant to counterparties - the borrower's risk, the worker's productivity - so privacy protects misrepresentation, degrades matching, and taxes the honest, who must pool with the concealed; on this view information flow is efficient and privacy regulation a rent to the low-quality. The rival tradition observed that information flows have *distributive* and *dynamic* consequences the efficiency story omits: disclosure enables the price discrimination of Section 2.2 (surplus moves, even when allocation improves), the classification situations of Chapter 8 (the score follows the person into markets it was not made for), and chilling effects on experimentation, association, and speech whose costs appear nowhere in the transaction that caused them. Acquisti et al.'s (2016) synthesis is the field's settled frame: privacy has no general welfare sign - whether information flow helps or harms depends on who uses it, for what, in which market structure - so the productive questions are contextual: which flows discipline sellers, which enable exploitation, and who holds the allocation of control. The reader will recognize the arc: it is Jones and Tonetti's (2020) allocation result from Section 6.7 arriving from the micro side.

The paradox that isn't, and the market that can't

The empirical literature's famous puzzle - people declare privacy priceless and sell it for a discount coupon - dissolves under the survey's analysis into structure rather than hypocrisy. Privacy preferences are real but *context-dependent* and *malleable*: valuations swing with framing, defaults, and salience, which is precisely why the interface design of Case 10.1's consent banners is where the battle is actually fought - the dark pattern is applied behavioural economics. And the market for privacy fails for textbook reasons the reader can now assemble: radical information asymmetry (the buyer of a 'free' service cannot observe the data flows she is pricing, the tracking asymmetry of Section 2.3 inverted); externalities (my disclosure trains the model that classifies you - the social graph makes privacy a public good, so individual consent under-protects even when informed); and the absence of the counterfactual (harms arrive later, elsewhere, unattributably - the attribution problem of Section 5.7 in reverse). These failures, not paternalism, are the economic case for the ex ante regulation of Chapter 10 - and the same failures predict regulation's characteristic disappointments, from consent theatre to the concentration effect of Johnson et al. (2023). Discussion questions: (1) Construct one market where more information flow is clearly efficient and one where it is clearly exploitative; identify the structural difference. (2) 'The privacy paradox is a measurement artefact plus a market failure.' Defend, using framing evidence and the externality argument. (3) If privacy is partly a public good, what follows for consent-based regimes - and which Chapter 10 instrument addresses the gap?

Applied Extension: Cybersecurity, Resilience and the Material Economy of Digitalization

The digital economy is frequently described as weightless, yet its most consequential failures reveal a material and security-intensive substrate. Cybersecurity is an economic problem because protection produces externalities: an insecure device, library, vendor or account can impose costs on parties that did not choose the vulnerability, while the benefits of precaution are difficult for buyers to observe before failure. This creates familiar information problems - credence qualities, moral hazard and underinvestment in prevention - now operating through software supply chains (Anderson & Moore, 2006).

Resilience adds a concentration dimension. Cloud regions, identity providers, content-delivery networks, payment rails and critical open-source components can become systemic single points of failure precisely because scale economies make centralization efficient in normal times. The economic question is therefore not whether concentration is efficient but

how much redundancy, interoperability and recovery capacity should be purchased against low-frequency, high-impact failure. The answer resembles insurance: resilience looks wasteful until the state of the world in which it is valuable occurs.

The same corrective applies to environmental analysis. Digital information is non-rival, but computation is not immaterial: data centres occupy land, consume electricity and water, rely on globally concentrated semiconductor supply chains, and generate hardware turnover and e-waste. The marginal cost of copying a file may approach zero while the marginal and fixed costs of producing computation remain economically and environmentally significant. AI makes this distinction particularly visible because training and inference transform electricity, chips and cooling capacity into explicit constraints on the supposedly weightless economy.

The size of this footprint is contested but no longer negligible. Freitag et al. (2021), reviewing competing estimates, place information and communication technology at roughly 2.1-3.9 per cent of global greenhouse-gas emissions once supply chains and device manufacturing are included, and they warn that efficiency gains have historically been outpaced by growth in use - the rebound effect familiar from energy economics. At the level of individual systems, Luccioni et al. (2024) show that the energy cost of serving a model varies by orders of magnitude across tasks, with general-purpose generative models consuming far more energy per query than task-specific ones; design choices at the application layer therefore have material consequences. Crawford (2021) extends the accounting beyond carbon to minerals, water and labour and frames artificial intelligence as an extractive industry. For the economist, these findings describe unpriced externalities. For the sociologist and the media scholar, they raise the questions of who bears the local costs of data centres and whose narratives of ‘weightless’ technology prevail.

Chapter 7. Digital Labour and the Transformation of Work

Learning objectives. After studying this chapter, the reader should be able to: (1) map the varieties of digital labour - local gig work, remote crowdwork and freelancing, and platform-adjacent conventional employment - and identify what platform mediation changes in each; (2) define algorithmic management and analyse it as a mode of labour control, comparing it with bureaucratic and Taylorist antecedents; (3) present the task-based economics of technological change, including routine-biased technical change, job polarization, and the displacement-reinstatement accounting; (4) evaluate the forecasting controversy over automation and employment; (5) render visible the hidden labour of the digital economy - moderation, microwork, and free labour - and explain why its invisibility is structural rather than accidental; and (6) connect the classification struggle over employment status to the platform economics of Chapters 3 and 5.

KEY CONCEPTS

gig work; crowdwork; algorithmic management; routine-biased technical change; polarization; displacement; reinstatement; invisible labour; remote work; worker classification

The preceding chapters analysed the digital economy from the standpoint of firms and markets; this chapter turns to work. The turn is not a change of subject but of perspective: the marketplace of Chapter 5 is, seen from below, a workplace without an employer of record; the algorithmic systems of Chapter 6 are, applied to workers, a management regime; and the platform-dependent entrepreneur of Chapter 4 is, in the sociology of work, a labour classification dispute walking. Vallas and Schor (2020), surveying the field, propose the orienting image this chapter adopts: the platform is neither simply a firm nor simply a market but a *permissive potentate* - it externalizes responsibility to workers while retaining, through code, the authority to structure their work - and the analytical task is to specify how that authority operates and what it costs and yields for those beneath it.

7.1. Gig Work, Crowdwork, and Algorithmic Management

Platform labour divides, geographically and analytically, into two families. *Local* gig work - ride-hailing, delivery, care and domestic tasks - is place-bound service labour matched by app; its economics are those of the on-demand marketplaces of Chapter 5, and its workforce is recruited by low entry barriers and schedule flexibility that survey after survey confirms workers genuinely value. *Remote* crowdwork and platform freelancing - microtasks, design,

programming, writing - detaches the labour market from geography entirely, creating what Graham and Anwar (2019) call a planetary labour market - its development consequences examined by Graham et al. (2017) - in which workers in lower-income countries reach global demand but also compete in a global oversupply that presses piece rates downward and shifts bargaining power decisively toward buyers. Wood et al. (2019) capture the resulting duality in the phrase 'good gig, bad gig': the same platform architecture yields real autonomy - over schedule, place, and task selection - alongside low pay, insecurity, overwork at the mercy of demand troughs, and the isolation of a workforce that never shares a canteen. The duality is not a contradiction in the evidence but the structure of the phenomenon: flexibility and precarity are jointly produced by the same contractual and technical design.

Form of labour	Market reach	Control mechanism	Primary vulnerability
Local gig work	Local	Dispatch, dynamic pricing, ratings	Unpaid waiting, classification, demand volatility
Remote freelancing	Global	Search ranking, reputation, client selection	Global oversupply, fee pressure, account dependency
Microwork / ghost work	Global and task-fragmented	Task acceptance, requester approval	Very low bargaining power, invisible unpaid time
Creator work	Global attention market	Recommendation, monetization eligibility	Income concentration, algorithmic volatility
Conventional digital work	Firm / hybrid	Managerial and algorithmic monitoring	Work intensification, surveillance, skill displacement

Table 7.1. Varieties of digital labour and the mechanisms through which platforms allocate control and risk.

That design's core is *algorithmic management*: the delegation to software of the functions - assignment, direction, evaluation, discipline - that management science assigns to supervisors. Rosenblat and Stark (2016), in the canonical case study of ride-hailing drivers, document its instruments: opaque dispatch and pricing, information asymmetries engineered into the interface (the destination concealed until acceptance), rating systems that convert customers into unpaid supervisors, and nudges - surge notifications, streak bonuses, 'you're almost there' prompts - that steer labour supply without commands. Kellogg et al. (2020) systematize the field: algorithms *direct* (recommending and restricting tasks), *evaluate* (rating and ranking), and *discipline* (rewarding, demoting, deactivating), and the ensemble constitutes a contested terrain of control descended from, but distinct from, Taylorism and bureaucracy - comprehensive like the former, rule-bound like the latter, but opaque, instantaneous, and

experienced as unappealable. The sociological punchline connects to Chapter 6: the black-box problem of algorithmic decision-making becomes, at work, a due-process problem - the deactivated driver confronts a decision with no stated reason, no named decider, and no procedure - and the classification struggle of Section 7.3's coda begins here, since the more granularly the algorithm directs the work, the harder it is to maintain that no employment relationship exists.

Algorithmic allocation does not abolish older inequalities; it can reproduce them through apparently neutral channels. Cook et al. (2021), analysing more than a million ride-hailing drivers in the United States, find a gender earnings gap of roughly 7 per cent even though the platform's pay formula is gender-blind. The gap is explained almost entirely by differences in experience on the platform, in where drivers choose to work and in driving speed - choices shaped by safety concerns, caring responsibilities and time constraints outside the algorithm. The finding cuts both ways. It refutes the claim that platform pay is discriminatory by design, and it shows that a formally neutral rule can still reward circumstances that are unequally distributed between men and women.

COMMON MISCONCEPTION

Automation should not be analysed only at the occupation level. Technologies usually change bundles of tasks; some are substituted, others complemented, and entirely new tasks can be created. Employment outcomes therefore depend on task composition and organizational redesign.

7.2. Skills, Polarization, and Technological Unemployment

Behind the platform foreground stands the longer question: what does digitalization do to the structure of employment as a whole? The modern economics of the question is *task-based*, and its foundation is Autor et al. (2003): computers substitute for labour in *routine* tasks - those reducible to explicit rules, whether manual (assembly) or cognitive (bookkeeping, filing) - and complement labour in non-routine abstract tasks (problem-solving, management, analysis), while neither substituting for nor substantially complementing the non-routine manual work (care, cleaning, driving, service) that demands situational adaptability. The prediction is not uniform displacement but *polarization*: growth at the top and bottom of the skill distribution and erosion in the routine middle - precisely the pattern Goos et al. (2014) document across sixteen European countries, attributing the lion's share to routine-biased technical change rather than to trade alone. Autor (2015) draws the synthetic lesson of two centuries: automation anxiety is perennially half-right - tasks are destroyed on schedule - and perennially half-wrong, because it reasons from substitution while forgetting complementarity,

income effects, and the new tasks that technology and rising demand create; the durable problem is not the quantity of work but its distribution, quality, and wages.

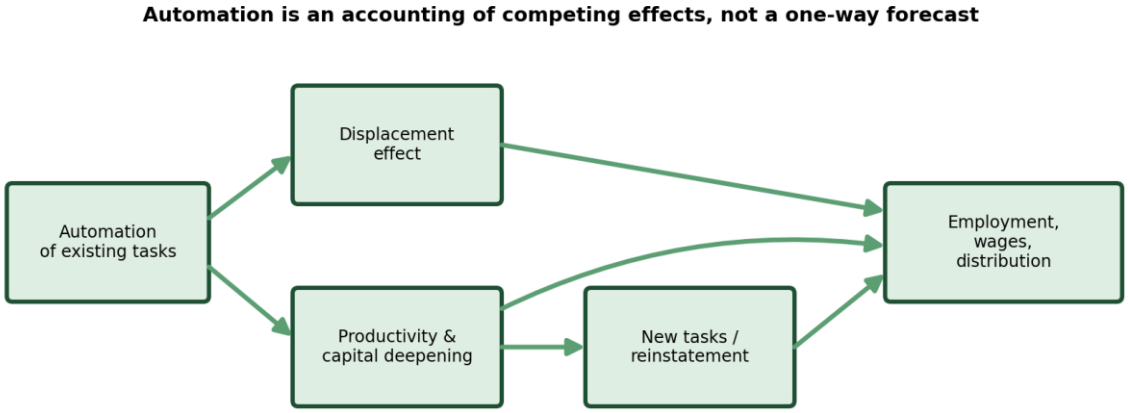

Figure 7.1. Displacement, productivity and reinstatement in task-based automation theory.

The forecasting controversy of the 2010s illustrates the stakes of method. Frey and Osborne (2017), estimating occupation-level 'computerisability', put forty-seven per cent of United States employment at high risk - a number that circled the globe and set the tone of a decade's automation discourse. The task-based rejoinder was that occupations are bundles: automating some tasks within a job more often transforms than eliminates it, and task-level re-estimates for OECD countries cut the high-risk share to roughly one job in ten to one in seven (Arntz et al., 2016; Nedelkoska & Quintini, 2018). Acemoglu and Restrepo's (2019) accounting, introduced in Chapter 6, absorbs the controversy into theory: displacement is real and measurable, but so are productivity effects and the reinstatement of labour in new tasks, and the balance is institutional - tax codes that subsidize capital relative to labour, research directions that pursue 'so-so automation' (displacing workers without large productivity gains), and education systems that lag task change all tilt the outcome. Generative AI, with its inverted exposure gradient (Eloundou et al., 2024), re-runs the debate one occupational floor higher; the reader now possesses the framework with which to meet it.

7.3. Invisible Labour: Moderation, Microwork, and Free Labour

The digital economy's self-presentation is one of automation; its production function tells otherwise. Gray and Suri (2019) name the gap *ghost work*: the human labour - tagging images, transcribing audio, verifying identities, cleaning training data, stepping in when models fail - that makes 'artificial' intelligence function, performed through microtask platforms by a

globally distributed, piece-paid, and contractually invisible workforce. The invisibility is structural, not incidental: the product is marketed as automated, so the humans in the loop must be architecturally concealed; the work is decomposed into seconds-long tasks, so no task is a job and no requester an employer; and the workforce is dispersed across jurisdictions, so no labour law finds its addressee. The AI systems of Chapter 6 thus rest on a supply chain of exactly the labour the automation discourse declares obsolete - a point with obvious consequences for how the productivity claims of that discourse should be audited.

A second stratum of hidden work is *content moderation*: the industrial-scale screening of user-generated material - violence, exploitation, abuse - that keeps commercial platforms advertiser-safe. Roberts (2019) documents its organization: outsourced, offshored, bound by non-disclosure, paid at business-process-outsourcing rates, and exposed to psychological harm that the industry treats as an externality; and her analysis makes the media-theoretical point that Chapter 9 will need - moderation is editorial labour, the human judgment inside the allegedly neutral platform, and its concealment sustains the fiction that no editing occurs. The third stratum is older than both: Terranova (2000), writing before the platform era, identified *free labour* - the voluntary, unwaged, pleasurable production of content, community, and attention by users themselves - as a constitutive feature of the digital economy, not a marginal one. Two decades of social media, reviews, open-source contribution, and the aspirational creator labour of Chapter 5 (Duffy, 2017) have made her essay look less like a provocation than a description: the audience commodity of Smythe (1977) now manufactures itself.

These strata converge on the chapter's closing question: the *classification struggle*. Whether a platform worker is an independent contractor or an employee determines minimum wage, social insurance, collective bargaining, and who bears demand risk - and, as Chapter 5 noted, the unit economics of several on-demand models were built on the contractor answer. De Stefano (2016) states the legal-theoretical stakes: platform work is 'just-in-time' labour whose casualization is enabled by the rhetoric of entrepreneurship, and the doctrinal tests of subordination were written for the foreman, not the algorithm - yet algorithmic management, once its direction, evaluation, and discipline are documented as in Section 7.1, satisfies the substance of subordination remarkably well. Jurisdictions have accordingly diverged - reclassification here, intermediate statuses there, presumptions and carve-outs elsewhere - and the struggle continues in courts and legislatures examined in Section 7.7 and Chapter 10. For the sociologist, the deeper observation anticipates Chapter 8 and the work of Fourcade and Healy (2017): classification is not the recognition of a pre-existing fact but the allocation of a

fate, and the digital economy has made the category of 'employee' - like the categories of creditworthy, relevant, and visible - an object of engineered contest.

WORKED EXAMPLE

Scenario. A crowdworker receives EUR 0.08 per accepted task. One task takes 70 seconds, but unpaid searching, qualification tests, and rejection risk add 30 seconds on average.

Calculation. Effective throughput is 36 tasks per hour. Gross effective hourly earnings are 36 × EUR 0.08 = EUR 2.88 before platform fees, taxes, equipment costs, and longer idle spells between available tasks.

Interpretation. The economically relevant wage is based on total labour time, not only the time visible inside the paid microtask.

THREE LENSES

Economics asks: How do technologies reallocate tasks and change productivity, wages, and bargaining power?

Sociology asks: How are control, precarity, identity, and inequality reorganized by digital management?

Media studies asks: How do interfaces, ratings, dashboards, and opaque metrics communicate managerial authority?

7.4. Chapter Summary and Discussion Questions

This chapter analysed work in the digital economy across three registers. Platform labour was mapped into local gig work and planetary remote work (Graham et al., 2017), its constitutive duality of autonomy and precarity documented (Wood et al., 2019), and algorithmic management specified as direction, evaluation, and discipline by code (Rosenblat & Stark, 2016; Kellogg et al., 2020), with Vallas and Schor's (2020) synthesis framing the platform as a structure that externalizes responsibility while retaining control. The aggregate economics of technological change was presented in task-based form - routine-biased technical change (Autor et al., 2003), polarization (Goos et al., 2014), the two-century verdict of Autor (2015), the forecasting controversy (Frey & Osborne, 2017), and the displacement-reinstatement accounting (Acemoglu & Restrepo, 2019). Finally, the hidden labour of the digital economy - ghost work (Gray & Suri, 2019), commercial content moderation (Roberts, 2019), and free labour (Terranova, 2000) - was rendered visible, and the classification struggle over employment status (De Stefano, 2016) was framed as the juridical face of everything the chapter described.

Discussion questions. (1) 'Flexibility and precarity are jointly produced.' Using Wood et al. (2019), explain how the same design features generate both, and propose one reform that preserves the former while reducing the latter. (2) Compare algorithmic management with

Taylorism along three dimensions: visibility of rules, granularity of monitoring, and avenues of contestation. Which regime concedes workers more, and why? (3) Reconstruct the Frey-Osborne estimate and the task-based critique. What, precisely, does each assume about the divisibility of occupations? (4) Why must ghost work be invisible? Connect the marketing of automation, the decomposition of tasks, and the geography of the workforce. (5) Is the content moderator a journalist? Use Roberts (2019) to argue the affirmative and the platform's terms of service to argue the negative, and state what turns on the answer for Chapter 9's concerns. (6) Draft the algorithmic-age test of employment subordination in one sentence, and apply it to a driver, a crowdworker, and a creator.

7.4a. In Focus: Digital Labour in the Global South

The gig work, crowdwork and "invisible labour" analysed in this chapter are globally distributed, and the distribution is not neutral. Much of the content moderation that keeps large platforms usable is performed by outsourced workers in Kenya, the Philippines and India, often on traumatic material for low pay; the data-labelling and "human-in-the-loop" work behind machine learning (Section 6.7) is routed through microwork platforms to workers in lower-income countries; and ride-hailing and delivery apps have become major urban employers across Africa, South Asia and Southeast Asia, exporting the algorithmic-management practices of Section 7.1 into very different labour-law regimes.

Attending to this geography complicates the Northern debate about "the future of work": for many workers the platform is not a supplement to a salaried job but the primary, precarious source of income, and the "free labour" that trains AI systems is often paid labour made invisible by distance and subcontracting. Seminar question: using this chapter's concepts, contrast the algorithmic management of a delivery rider in a high-income city with one in a low-income city - what protections differ, and why?

7.4b. Numerical and Data Exercises

Exercise 7.A (effective pay and earnings floors). A driver is logged in for 50 hours in a week and engaged (carrying passengers) for 32 of them. Gross fares after the platform's commission are EUR 650; vehicle costs are EUR 0.22 per kilometre over 1,250 km. (a) Compute net earnings and net pay per logged-in hour and per engaged hour. (b) An earnings floor guarantees gross fares of at least 120% of a EUR 12 hourly minimum wage for engaged time only. Is it

binding? (c) Would it bind if it applied to logged-in time? Why does the distinction matter (Case 7.1)?

Exercise 7.B (exposure versus displacement). An occupation's working time is divided among three tasks: A (30%), B (50%) and C (20%). A generative AI tool halves the time needed for task A and cuts the time for task C by a quarter; task B is unaffected. (a) What share of working time is exposed? (b) By how much does total working time for the same output fall? (c) Why does neither number tell you whether employment in the occupation will fall (Section 6.4)?

Indicative answers. 7.A (a) Costs = 0.22 × 1,250 = EUR 275; net = EUR 375; EUR 7.50 per logged-in hour and about EUR 11.72 per engaged hour. (b) Floor = 1.2 × 12 × 32 = EUR 460.80; gross fares of EUR 650 exceed it, so it does not bind. (c) On logged-in time the floor would be 1.2 × 12 × 50 = EUR 720, so the platform would owe a top-up of EUR 70. Counting only engaged time leaves the cost of waiting - which the dispatch algorithm itself shapes - with the worker. 7.B (a) Tasks A and C, i.e. 50% of working time. (b) 0.30 × 0.5 + 0.20 × 0.25 = 0.20, a 20% reduction. (c) Employment depends on how demand responds to lower costs, on whether freed time is redeployed to other or new tasks, and on institutions - the displacement, productivity and reinstatement effects of Acemoglu and Restrepo (2019).

7.5. Case Studies and Debates

The materials below follow the chapter's arguments into a courtroom-and-ballot-box struggle over classification, a survey of the microwork economy's actual conditions, and the era's most sweeping proposed remedy.

Case 7.1. California's classification war: AB5 and Proposition 22

Nowhere did the classification struggle of Section 7.3 escalate more completely than in California, home jurisdiction of the ride-hailing industry. In 2019 the state legislature enacted Assembly Bill 5, codifying a strict presumption - the 'ABC test' - under which a worker is an employee unless the hirer proves the work is free from its control, outside its usual course of business, and part of the worker's independent trade. The second prong was aimed at the heart of the platform position: driving passengers is not plausibly outside a ride-hailing firm's usual course of business, so the statute threatened wholesale reclassification, with the minimum-wage, insurance, and bargaining consequences of Section 7.3. The platforms' response was a lesson in the political economy of code and capital jointly deployed: they declined to reclassify

while litigating, and financed Proposition 22 - at the time the most expensive ballot-initiative campaign in American history - which voters approved in 2020. The initiative exempted app-based drivers from AB5 and substituted a bespoke third status: independent contractors with limited benefits (a healthcare stipend, an earnings floor covering *engaged* time only) and, crucially, a supermajority requirement entrenching the settlement against legislative revision. Years of constitutional litigation followed, with the initiative ultimately surviving.

Dubal (2017) supplies the frame that makes the episode legible: the century-long contest over whom the law calls an employee has always been a struggle over identities as much as costs - the same driver can be narrated as wage worker or as micro-entrepreneur, and Section 4.4's founder mythology was deployed, in the campaign, to recruit drivers' own self-understanding to the contractor side. The economics are equally instructive: the earnings floor covering only engaged time prices the waiting that Section 7.1 showed the algorithm itself induces, and analysts disputed whether the guaranteed floor, net of unpaid waiting and expenses, exceeded or fell short of the employment counterfactual. Seminar questions: (1) Apply the ABC test and the algorithmic-management evidence of Section 7.1 to a delivery platform; which prong decides the case? (2) Was Proposition 22's third status an innovation (portable benefits fitted to flexible work) or an evasion (employment minus its price)? Specify the empirical facts on which your answer turns. (3) The platforms bought a change in the law's categories. Analyse the episode with Baumol (1990): toward what did the reward structure steer entrepreneurial talent?

Case 7.2. The microwork economy: conditions on the invisible shop floor

What are the actual terms of the ghost work of Section 7.3? The International Labour Organization's survey of thousands of crowdworkers across five major microtask platforms (Berg et al., 2018) remains the most systematic answer, and its findings give the chapter's abstractions their numbers. Averaged across platforms and regions, earnings ran to a few dollars per hour - below local minimum wages for a substantial share of workers, especially once *unpaid* labour is counted: roughly twenty minutes of unpaid work accompanied every hour of paid work, spent searching for tasks, taking unpaid qualification tests, and vetting requesters, the search costs of Chapter 2 shifted onto the worker. Rejection of completed work without pay or explanation - the requester's unilateral privilege on most platforms - was widespread, and with it the due-process vacuum of Section 7.1; a large minority of respondents relied on microwork as their primary income, against the industry framing of pocket-money supplementation; and the workforce was strikingly educated, with majorities holding post-

secondary qualifications - the planetary labour market of Graham et al. (2017) arbitraging credentialed labour at piece rates.

The case also documents the countervailing institutions forming in the vacuum: worker forums and requester-rating tools through which crowdworkers share information about payers - reputation systems (Chapter 8) built bottom-up, pointing the asymmetry the other way - and the ILO's own proposal repertoire, from platform codes to the extension of wage floors across contractual form. Seminar questions: (1) Compute the effective hourly wage of a worker paid per task once search, qualification, and rejection are counted; which platform design choices move each component? (2) Requesters can reject work without pay; buyers on goods marketplaces cannot take delivery and refuse payment. What explains the difference, and what would eliminate it? (3) Design a portable, cross-platform reputation for crowdworkers (Chapter 8); who would resist it, and why?

Debate 7.1. Universal basic income: remedy or resignation?

Motion: a universal basic income (UBI) is the appropriate institutional response to digital automation. The affirmative case assembles the book's own materials: if the automation accounting of Acemoglu and Restrepo (2019) can tip toward displacement, if platform work strips the benefits historically attached to employment (Section 7.3), and if the data and free labour of Chapters 6 and 7 constitute uncompensated contribution to the digital economy's surplus, then an unconditional floor - paid to all, free of means tests and work requirements - secures subsistence independently of a labour market growing less reliable as a distribution mechanism. It would also restore bargaining power (the credible exit from bad gigs that Section 7.1's workers lack) and price, however roughly, the collective inheritance of technology (Van Parijs & Vanderborght, 2017). *Against the motion:* the diagnosis is premature and the prescription mistargeted - Autor's (2015) two centuries counsel that work persists while its quality and distribution deteriorate, problems addressed by wage subsidies, portable benefits, and labour standards at a fraction of UBI's fiscal mass, which is enormous precisely because universality pays the majority who need nothing. Decoupling income from contribution risks eroding both the public legitimacy of redistribution and the reinstatement channel - institutions should steer technology toward augmenting labour (Acemoglu & Johnson, 2023), not compensate for surrendering it. Finally, the empirical record of pilots shows well-being gains but little evidence on the macro-fiscal and labour-supply questions on which the policy actually turns. The debate should require each side to specify financing, the fate of existing social

insurance, and - the hinge - whether they believe the displacement scenario or merely insure against it.

7.6. In Depth: Remote Work and the New Geography of Labour

The platform labour of Sections 7.1 and 7.3 is the digital transformation of work at its most visible; this advanced section treats its quieter and quantitatively larger counterpart - the untethering of conventional employment from the workplace - whose economics were established by field experiment before a pandemic ran the experiment on the world.

The experiment before the experiment

The touchstone study predates the mass remote era by half a decade and owes its authority to design. Bloom et al. (2015) randomized working from home among volunteer call-centre employees of a large Chinese travel firm: treated workers performed at home four days a week for nine months against an office control group, with output, hours, and attrition measured administratively. The results were unambiguous within their domain: home workers' performance rose by roughly thirteen per cent - more minutes worked per shift (fewer breaks and sick days, no commute leakage) and more calls per minute (a quieter environment) - while quits fell by half and reported satisfaction rose. Two further findings proved as consequential as the headline. Promotion rates of home workers fell, conditional on performance - early evidence of a visibility penalty that would become the hybrid era's central anxiety - and when the firm rolled the option out and let workers re-choose, half the volunteers returned to the office, and selection did substantial work: those who chose home were those who performed well there, so the *option's* value exceeded the *mandate's*. The study's discipline matters as much as its findings: it measured routinized, individually metered work, and its authors resisted extrapolation to collaborative and creative tasks - a restraint the subsequent debate did not always imitate.

The pandemic reallocation and its settlement

The involuntary global experiment of 2020 converted a marginal arrangement into a mass one and, as the emergency receded, settled into the pattern documented across the survey and mobility evidence synthesized by Barrero et al. (2023): work from home stabilized at a multiple of its pre-pandemic level, overwhelmingly in the *hybrid* form for professional employment, with full remote concentrated in a minority of jobs and firms. The economics of the settlement are legible with the book's tools. Workers value the arrangement concretely - the commute time returned is the largest in-kind pay rise in decades, and survey respondents price hybrid access at a meaningful share of salary - so remote amenity entered wage bargaining as compensating differential: firms offering it capture labour at a discount, a margin that quietly

restrained measured wage growth in the reallocation years. The technology stack that made it viable - the cloud collaboration and video infrastructure of Section 5.3 - had been built for other purposes and was waiting. The productivity evidence settled into a sober middle (modest positive effects for hybrid, more contested for full remote, with coordination and mentoring the recurring costs). The distributional facts attached immediately: remote-capable work is concentrated among the educated and well paid, so the amenity accrued along exactly the gradients of Section 8.2, while the manual-interactive occupations of Section 7.2 - care, logistics, food, the pandemic's 'essential' work - were untethered from nothing.

Second-order geographies

The durable consequences are spatial, and they knit the section to the rest of the book. Within metropolitan areas, the reallocation repriced real estate along the commuting gradient - central office districts lost the daily population that sustained their service economies, residential demand shifted outward, and the municipal fiscal model built on commercial property values entered a slow crisis the reader should recognize as an infrastructure externality of a labour-market change. Between regions, remote work loosened, without severing, the tie between where value is produced and where producers live: satellite towns and amenity regions gained high-income residents whose employers remained elsewhere, importing income and gentrification pressures together. And between countries, the frontier case is the one Section 7.1 already met: if the job leaves the building, some jobs can leave the country, and the planetary labour market of Graham et al. (2017) extends upward from microtasks toward professional services - constrained less by technology than by the institutions this book keeps finding at the bottom of things: time zones and trust (Chapter 8), employment law and taxation (Section 7.3), and the immigration rules that remote hiring quietly arbitrages. Discussion questions: (1) Reconstruct the Bloom et al. (2015) design and explain what the re-choice phase reveals that the randomization alone could not. (2) Model hybrid access as a compensating differential; who ultimately captures its value - workers, shareholders, or landlords along the new gradient? (3) The visibility penalty is a measurement problem: proximity proxies for contribution. Connect it to the metrics of Section 8.3 and design the evaluation regime that removes it - then state its new failure mode. (4) Which professional services will cross borders next, and which of the constraining institutions binds first?

7.7. In Depth: Regulating Platform Work - Europe's Presumption of Employment

Sections 7.1 and 7.3 described algorithmic management and the classification struggle, and Case 7.1 followed that struggle through California's courts and ballot boxes. This advanced section examines the European answer, which is now becoming law across the Union, because it tests whether the mechanisms identified in this chapter can be governed by statute.

From courts to a directive

European jurisdictions first moved case by case. The United Kingdom's Supreme Court held that ride-hailing drivers were 'workers' entitled to the minimum wage and paid holiday, reasoning from the platform's control over pricing, contract terms and performance ratings (Uber BV v Aslam, 2021). Spain went further by statute: its 2021 'riders' law' introduced a presumption that delivery riders managed through a digital platform are employees, and it required firms to disclose to worker representatives the parameters of algorithms that affect working conditions (Ley 12/2021, 2021). The European Union generalized both approaches in the Platform Work Directive, which member states must transpose by 2 December 2026 (European Parliament & Council of the European Union, 2024a).

Two pillars: presumption and algorithmic management

The directive rests on two pillars. The first is a rebuttable legal presumption of an employment relationship where the facts indicate control and direction. The burden of proving genuine self-employment shifts to the platform, which converts the evidentiary asymmetry of Section 7.1 - the platform holds the data, the worker does not - into a legal rule. The second pillar applies to everyone working through platforms, including the genuinely self-employed. Platforms must inform workers about automated monitoring and decision-making systems and ensure human oversight of significant automated decisions. Workers gain the right to an explanation and to review of decisions such as account suspension, and platforms may not process certain categories of personal data, among them data on emotional or psychological states and on private conversations. In the vocabulary of this chapter, the first pillar addresses classification and the second addresses the due-process vacuum of algorithmic management.

What to watch

The directive leaves the facts that trigger the presumption largely to national law, so transposition is likely to reproduce, within Europe, the divergence familiar from Case 7.1, and

implementation has so far been uneven across member states. The economic questions are those raised in Chapter 5: whether on-demand models remain viable once employment costs are internalized, whether platforms respond by reducing flexibility or by redesigning control so as to avoid the presumption, and how prices, service coverage and the number of workers adjust. Discussion questions: (1) Compare the ABC test of Case 7.1, the UK 'worker' category and the EU presumption: which allocates the burden of proof most effectively, and why does the burden matter? (2) The algorithmic-management rules also apply to the self-employed. Which of the functions identified by Kellogg et al. (2020) - direction, evaluation, discipline - do they constrain most? (3) Predict one strategy by which a platform might redesign its control system to avoid the presumption, and describe the evidence a labour inspector would need to see through it.

Chapter 8. The Sociology of Digital Markets: Trust, Networks, and Inequality

Learning objectives. After studying this chapter, the reader should be able to: (1) apply the embeddedness perspective to online exchange and explain how digital markets engineer the trust they cannot presuppose; (2) present the economics and the pathologies of reputation and feedback systems, including the empirical evidence on their information content and their discriminatory side-effects; (3) distinguish the successive levels of the digital divide - access, skills and usage, and outcomes - and relate them to the reproduction of social inequality; (4) analyse quantification and algorithmic classification as social processes that create new forms of stratification; and (5) articulate what a sociological analysis adds to the economic account of digital markets developed in Chapters 2 through 5.

KEY CONCEPTS

embeddedness; trust; reputation; verification; digital divide; skills; usage inequality; outcome inequality; quantification; classification; social capital; institutional trust

The economic chapters of this book repeatedly reached conclusions that pointed beyond economics: verification costs are solved by reputation (Chapter 2), whose social life economics does not theorize; platform ignition requires trust among strangers (Chapter 3), which markets presuppose rather than produce; entrepreneurial opportunity travels through networks (Chapter 4), which are socially closed; and algorithmic classification allocates fates (Chapters 6 and 7), which is stratification by another name. This chapter supplies the missing theory. Its charter is

Granovetter's (1985) embeddedness thesis, introduced in Chapter 1: economic action is embedded in concrete networks of social relations, and the atomized market of theory is achieved - where it is achieved - by institutional work. The digital economy is the largest experiment ever run on that thesis: an attempt to conduct exchange among strangers at planetary scale by *engineering* the relational preconditions - trust, reputation, identity - that face-to-face markets inherited from social life. The experiment's results, this chapter argues, vindicate the sociology twice over: the engineering works, which shows the preconditions were real; and the engineering has social consequences its designers did not intend, which shows that institutions are never merely technical.

8.1. Embeddedness and Trust in Online Exchange

The founding puzzle of electronic commerce was stated early and precisely: why would anyone send money to a pseudonymous stranger for unseen goods? Resnick and Zeckhauser (2002), analysing the feedback system of the pioneering auction marketplace, gave the canonical answer: reputation mechanisms convert the shadow of the future into a market institution. Where repeated interaction between the same partners is unlikely, the platform pools every partner's experience into a public score, so that each seller effectively plays a repeated game against the market as a whole; the expectation of future sales disciplines present honesty. Their empirical findings set the field's agenda as much as their theory: feedback was overwhelmingly, almost implausibly positive - negative ratings were rare enough to suggest under-reporting through fear of retaliation and simple non-return - and yet the system worked in the aggregate: it was, in their phrase, sufficient to keep the market honest even if noisy in each instance. Trust, in short, could be manufactured - imperfectly, but at scale.

Tadelis (2016) synthesizes the two subsequent decades of economics on these systems and confirms both halves of the founding result. Reputation carries real information - better scores command price premia and survival advantages, and experimental manipulations show buyers respond causally to ratings. Yet reputation systems also exhibit systematic distortions: grade inflation that compresses scores toward the top until the signal lives in the second decimal; reciprocity and retaliation in two-sided rating, which platforms learned to suppress with simultaneous-reveal designs; selection into reviewing that over-represents extremes; and outright manipulation through purchased or fraudulent reviews, an adversarial arms race that platforms police as a core cost of business. The design details - who may rate whom, when scores are revealed, how histories decay - are therefore not technicalities but the constitutional law of these markets, which is the embeddedness thesis restated in engineering form: the market

works because a designed social institution does its relational work. What the economics of reputation understates, the sociology supplies: the rated do not merely receive scores; they live under them - the emotional labour of the five-star service encounter, the driver's anxiety documented in Chapter 7, the strategic self-presentation that ratings compel - so that the reputation system is simultaneously an information mechanism and a discipline, Foucault's observation arriving by way of customer satisfaction surveys.

The discipline, moreover, is not neutral among persons. Edelman et al. (2017), in a field experiment on a short-term accommodation platform, found that guest applications from distinctively African-American names were accepted roughly sixteen per cent less often than identical applications under white-sounding names - discrimination enabled, ironically, by the very identity disclosure (names, photographs) that the platform introduced to build interpersonal trust. The finding generalizes into a design dilemma constitutive of digital markets: the information that manufactures trust also transmits the categories on which prejudice operates, so that platforms choose, whether they acknowledge it or not, a point on the trade-off between anonymity (less discrimination, less trust) and identity (more trust, more discrimination) - and subsequent design reforms (delaying photographs, instant booking) are movements along it. The episode is this chapter's thesis in miniature: an institution engineered for an economic function - verification, Chapter 2 - imports the social structure of its users, and the import is a design variable for which the platform bears responsibility.

COMMON MISCONCEPTION

The digital divide is not solved when access becomes widespread. Inequality can migrate from connection to device quality, skills, autonomy of use, platform literacy, safety, visibility, and the ability to convert digital participation into economic or social outcomes.

8.2. Digital Divides and Stratified Participation

Participation in digital markets presupposes participation in digitality itself, and that participation is stratified. The first generation of digital-divide research counted connections: who has access to computers and the Internet, along the predictable gradients of income, education, geography, and age - the binary divide that infrastructure expansion has narrowed without closing (Norris, 2001). The field's decisive turn was Hargittai's (2002) demonstration of a *second-level* divide: among the connected, skills differ systematically - in her experiments, the time and success with which users could complete basic information tasks varied enormously and along familiar social lines - so that formally equal access conceals substantively unequal capacity. DiMaggio et al. (2001), in the agenda-setting review cited in

Chapter 1, generalized the point into the research programme the field still follows: as access diffuses, inequality migrates into the *conditions* of use - autonomy of access, skill, social support, and the purposes to which the medium is put.

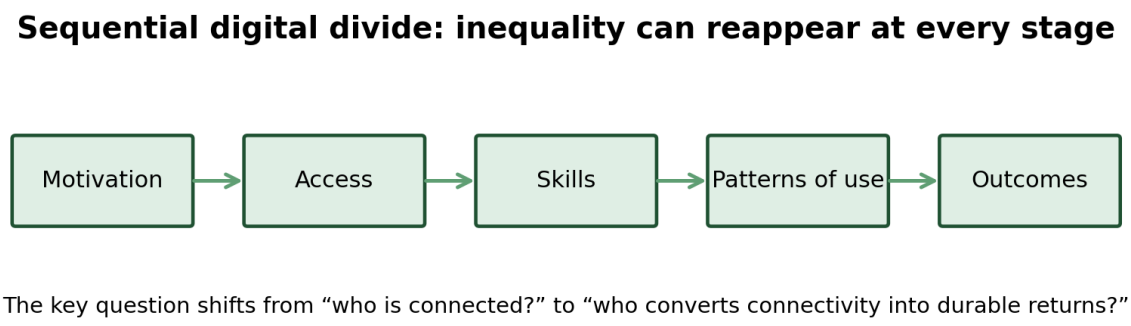

Figure 8.1. Sequential model of digital inequality: from motivation and access to outcomes.

Van Dijk (2020) organizes the accumulated findings into the sequential model that contemporary scholarship employs: motivation and attitudes condition physical access; access conditions the acquisition of skills - operational, informational, strategic; skills condition usage patterns; and usage patterns condition *outcomes* - the third-level divide: who converts connectivity into economic, social, and political returns. The model’s central empirical finding deserves the emphasis he gives it: at each stage, the already-advantaged gain more - higher-status users deploy the medium for capital-enhancing purposes (information, finance, career, civic voice) while lower-status use skews toward entertainment and consumption - so that the Internet, far from the equalizer of first-wave prophecy, operates by default as an *amplifier* of existing stratification. The relevance to this book’s earlier chapters is direct: the entrepreneurial opportunity of Chapter 4, the creator income of Chapter 5, and the remote work of Chapter 7 are all returns to digital participation, and the divide literature explains who is positioned to collect them. It also completes the measurement lesson of Chapter 1: aggregate adoption statistics, like aggregate GDP, conceal exactly the distribution in which the social question resides.

8.3. Reputation Systems and the Quantification of Social Life

The chapter’s final movement widens from markets to the society they are quantifying. The sociological theory of *commensuration* - Espeland and Stevens (1998) - supplies the frame: the transformation of qualities into a common metric is a social process with its own politics; metrics do not merely describe, they *discipline* - attention flows to what is measured, the

measured adapt to the measure (reactivity), and alternatives incommensurable with the metric become invisible. The digital economy is commensuration industrialized: the ratings of Section 8.1, the engagement metrics of Chapter 5, the risk scores of Chapter 6, and the worker ratings of Chapter 7 are one apparatus viewed at different points, converting conduct into comparable, rankable, decidable numbers. Mau (2019) names the aggregate result the *metric society*: a culture in which quantified reputation - scores, rankings, badges, followers - becomes a general currency of worth, sustaining permanent competition and self-optimization, and installing what Chapter 7 met as algorithmic discipline into everyday self-relation. Fourcade and Healy (2017) draw the stratification consequence that gives this chapter its destination. Modern institutions, they argue, increasingly 'see like a market': they classify persons through data traces into *classification situations* - credit tiers, risk categories, customer values, visibility ranks - that carry differential treatment; and the accumulated digital record functions as a new asset, which they name *übercapital*: a form of capital arising from one's position in the classificatory order, convertible into price, access, and opportunity. The scheme knits the book's threads into a single sociological claim: the data capital of Chapter 6, extracted from users, returns to them as classification; the classifications allocate the market chances that Weber called life-chances; and stratification acquires a new mechanism - neither property nor credential but *score* - that operates continuously, opaquely, and with the moralized aura of earned desert. Where Chapter 6 asked whether the classifications are biased, the sociological question is prior and remains open even when they are accurate: what kind of social order results when ever more of life is lived under permanent, consequential rating - and who audits the raters? The political-economy answer, such as it exists, is Chapter 10's subject; the media through which the metric society narrates itself are Chapter 9's.

Ethnographic evidence shows how the rated respond. In a study of an online labour platform, Rahman (2021) describes an 'invisible cage': because the platform's evaluation algorithm is opaque and changes without notice, freelancers experiment anxiously with their own behaviour - limiting whom they work with, avoiding risky projects and managing clients to protect their scores. Opacity, on this account, is not an incidental property of the system but an instrument of control, which links the sociology of quantification in this section to the due-process concerns of Section 7.1.

WORKED EXAMPLE

Scenario. Seller A has 100 ratings: 95 positive and 5 negative. Seller B has 20 ratings, all positive.

Reasoning. A naive positivity ratio prefers Seller B (100%) to Seller A (95%), although Seller A has far more evidence. A Bayesian or confidence-adjusted reputation score would discount the uncertainty attached to the small sample.

Interpretation. Reputation systems do not merely reveal trustworthiness; their statistical and interface design constructs the signal on which market access depends.

THREE LENSES

Economics asks: How does reputation reduce information asymmetry and transaction costs?

Sociology asks: Whose past behaviour becomes legible, how are categories constructed, and how do inequalities reproduce themselves?

Media studies asks: How do scores, badges, rankings, and profiles make reputational judgments visible and actionable?

8.4. Chapter Summary and Discussion Questions

This chapter supplied the sociological theory that the economic analysis of digital markets requires. Online exchange was shown to rest on engineered trust: reputation systems convert pooled experience into market discipline (Resnick & Zeckhauser, 2002), carry real but distorted information (Tadelis, 2016), subject the rated to continuous discipline, and transmit discrimination through the very identity cues that build trust (Edelman et al., 2017). Digital inequality was traced through its levels - access (Norris, 2001), skills (Hargittai, 2002; DiMaggio et al., 2001), and outcomes (van Dijk, 2020) - with the amplification thesis as the field's central result. Finally, quantification was theorized as commensuration (Espeland & Stevens, 1998), generalized as the metric society (Mau, 2019), and given stratification teeth through classification situations and übercapital (Fourcade & Healy, 2017), completing the argument that digital markets are social institutions whose design allocates not only efficiency but life-chances.

Discussion questions. (1) 'Reputation systems prove Granovetter right twice.' Explicate: what does their necessity show, and what do their side-effects show? (2) Using Tadelis (2016), diagnose why ratings inflate, and design a feedback mechanism that resists inflation without deterring participation. State the trade-offs you accept. (3) The Edelman et al. (2017) experiment identifies a trust-discrimination trade-off. Where should a platform sit on it, who should decide, and under what accountability? (4) Apply van Dijk's sequential model to one digital opportunity from this book - creator income, remote freelancing, or e-commerce

entrepreneurship - and identify the stage at which stratification bites hardest. (5) Give one example each of a metric that describes, a metric that disciplines, and a metric that has become übercapital; defend the classification. (6) ‘Accurate classification can still be unjust.’ Construct the strongest version of this claim using Fourcade and Healy (2017), and the strongest rebuttal using the economics of information from Chapter 2.

8.4a. Numerical and Data Exercises

Exercise 8.A (ranking under uncertainty). Seller A has 95 positive ratings out of 100; Seller B has 20 out of 20. (a) Compute each seller’s raw share of positive ratings. (b) Compute the lower bound of the 95% Wilson confidence interval, $[p + z^2/2n - z\sqrt{(p(1-p)/n + z^2/4n^2)}] / (1 + z^2/n)$, with $z = 1.96$. (c) Which rule should a platform display, and what does your choice do to newcomers (Section 8.6)?

Exercise 8.B (where the divide bites). In a population, 90% of high-income and 75% of low-income adults are online. Among those online, 60% of high-income and 30% of low-income users use the Internet for capital-enhancing purposes (job search, finance, learning). (a) Compute the ratio of access rates and the ratio of the shares of each whole group that use the Internet in capital-enhancing ways. (b) Which stage of van Dijk’s sequential model contributes more to the final gap? (c) Propose one policy aimed at each stage.

Indicative answers. 8.A (a) A: 95%; B: 100%. (b) A: about 0.888; B: about 0.839, so A ranks first once uncertainty is taken into account. (c) Confidence-adjusted scores protect buyers from small-sample noise but penalize newcomers with short histories; platforms often combine them with probationary badges or guarantees. 8.B (a) Access ratio $0.90 \div 0.75 = 1.2$; outcome shares 54% and 22.5%, a ratio of 2.4. (b) The ratio decomposes as 1.2×2.0, so the usage stage doubles the gap and contributes more than access. (c) For example, subsidized connectivity for access, and skills training linked to job search for usage.

8.5. Case Studies and Debates

The materials below extend the chapter's arguments through a design experiment on discrimination, the most ambitious state deployment of reputation infrastructure yet attempted, and the standing dilemma between anonymity and identity.

Case 8.1. Engineering against discrimination: what design can and cannot fix

The Edelman et al. (2017) experiment of Section 8.1 established that identity disclosure on accommodation platforms transmits racial discrimination; the natural follow-up question - can design remove it without removing trust? - has since acquired an experimental answer. Cui et al. (2020), running field experiments on the same market, found that the presence of positive *reviews* on a guest's profile essentially eliminated the acceptance-rate gap between distinctively African-American and white guest identities: where hosts could observe a track record, the statistical inference from name to risk lost its work, and behaviour converged. The finding is theoretically precise: it locates the documented discrimination substantially in *statistical* inference under information scarcity rather than in pure animus - hosts discriminated most when they knew least - and it identifies reputation, the chapter's central institution, as a substitute for the crude category. The platform's own reforms moved along the same frontier: delaying the display of names and photographs until after booking decisions, expanding instant booking that removes host screening altogether, and shrinking profile photographs.

The case's discipline lies in its limits. Reputation cures the inference only for those who possess it: the cold-start newcomer - disproportionately, by the divide logic of Section 8.2, the marginal entrant - still faces the categorical judgment, so the remedy reproduces a seniority gradient even as it flattens a racial one; review-based trust re-imports the inflation and reciprocity pathologies of Section 8.1; and animus-driven discrimination, by construction, survives information. Seminar questions: (1) Distinguish taste-based and statistical discrimination formally, and explain why reviews should cure only the latter; what observed pattern would indicate residual animus? (2) Rank the platform's three design reforms by predicted effectiveness and by cost to trust; justify the ranking with Sections 8.1 and 2.3. (3) Should anti-discrimination law mandate such designs? Contrast the regulatory logic of outcome auditing with design mandates, anticipating Chapter 10.

Case 8.2. The state as rating agency: China's social credit experiment

Section 8.3's metric society has a limiting case: the deliberate construction, by a state, of reputation infrastructure spanning economic and civic life. China's social credit initiative,

announced in a 2014 planning document, is routinely described abroad as a single, unified citizen score. The scholarly record describes something less cinematic and analytically more interesting - a heterogeneous ecology of instruments: regulatory blacklists and 'redlists' (most consequentially, a court-administered defaulters' list whose members face restrictions on flights, high-speed rail, and conspicuous consumption); municipal pilot schemes, some with point scores, most marginal in practice; and commercial credit-scoring products from platform firms, adjacent to but distinct from the state systems. Kostka's (2019) survey evidence supplies the case's most unsettling finding for Western intuitions: public approval of social credit systems among Chinese respondents was high, and *highest* among the educated and urban - interpreted by respondents less as surveillance than as an institutional substitute for the deficient interpersonal and commercial trust that Sections 8.1 and 2.3 identified as the precondition of market exchange.

The case therefore belongs to this chapter's theory, not merely to its politics. Where the platforms of Section 8.1 built reputation to enable transactions within a marketplace, the social credit project generalizes the ambition: commensuration (Espeland & Stevens, 1998) as state infrastructure, classification situations (Fourcade & Healy, 2017) issued by public authority, and the metric society (Mau, 2019) pursued as policy - with the crucial institutional difference that exit, the discipline that constrains a platform's scoring (users can leave), does not constrain a state's. Seminar questions: (1) Compare the defaulters' blacklist with a private credit score along four dimensions: consequences attached, contestation procedure, exit, and visibility. Which differences are intrinsic to statehood? (2) Use Kostka (2019) to construct the strongest *functional* defence of state reputation infrastructure in a low-trust market economy - then identify the premise a liberal critic must reject. (3) Western observers often analysed social credit as an exotic other; inventory the scoring systems - credit, tenant, fraud, risk - already governing life-chances in your own jurisdiction, and state precisely what distinguishes them.

Debate 8.1. Anonymity or identity: what should the network know?

Motion: digital platforms should be required to verify the real identity of their users. *For the motion:* pseudonymity is the subsidy on which the pathologies of three chapters run - the fraudulent reviews of Section 8.1, the harassment that drives the moderation labour of Chapter 7, the coordinated inauthentic behaviour of Chapter 9's information disorder. The trust machinery this chapter described is a costly workaround for an identity layer the network lacks. Verification aligns online accountability with the offline norm that consequential speech and commerce occur under one's name, and the divide literature itself (Section 8.2) shows that the

harms of anonymity fall heaviest on those least able to absorb them. *Against the motion:* the Edelman-Luca-Svirsky result is a standing demonstration that identity transmits discrimination - verification would generalize it across every market and forum. Anonymity is the protective infrastructure of whistle-blowers, dissidents, abuse survivors, and the merely unpopular, and its removal is regressive precisely along Section 8.2's gradients, since the privileged can afford their names. Verified-identity regimes create honeypots of exactly the credential data whose breach economics Chapter 6 described; and the empirical record of real-name policies shows modest civility gains against these costs. The debate's productive form is not binary: participants should design an identity *architecture* - which contexts require verification, which pseudonymity, which anonymity; who holds the mapping; and under what process it is pierced - and defend the allocation with the chapter's trust-discrimination trade-off stated explicitly.

8.6. In Depth: Designing Reputation - A Mechanism Designer's Notebook

Sections 8.1 and 8.3 analysed reputation systems as institutions; this advanced section takes the designer's chair, because every pathology Tadelis (2016) catalogued corresponds to a design lever, and working the levers teaches the deepest lesson of this chapter - that these systems are constitutions, and constitutional choices have no neutral setting.

What to elicit, and from whom

The first design layer is elicitation. Voluntary reviewing selects for extremes - the delighted and the furious report; the satisfied middle abstains - so the observed distribution misrepresents the experienced one, and platforms answer with solicited or verified-transaction reviews, trading selection bias for response fatigue and for the loss of the unverified customer's voice. Bilateral markets add the strategic layer Resnick and Zeckhauser (2002) first flagged: when both sides rate and ratings are sequential, fear of retaliation censors negative reports - the mechanism behind the near-uniform positivity of Section 8.1 - and the standard repair is the *simultaneous reveal*: both ratings are collected blind and published together, removing the retaliation channel at the cost of removing, too, the corrective dialogue that sequential feedback allowed. A subtler elicitation choice concerns *dimensionality*: the single star collapses cleanliness, accuracy, communication, and value into one number - maximal comparability, minimal information - while multi-dimensional ratings inform better and inflate differently, since raters spare the person on the overall score and confess on the sub-scores; platforms that mine sub-scores while displaying the headline number are performing, quietly, the commensuration politics of Espeland and Stevens (1998).

Aggregation, decay, and the constitution of second chances

The second layer is aggregation, and its choices are jurisprudence. Lifetime averages entrench history - the early stumble follows the seller forever - while windowed or decayed scores institutionalize forgetting: the half-life of a bad review is, precisely, the platform's statute of limitations, and choosing it allocates the trade-off between incentive maintenance (mistakes must cost) and rehabilitation (the cost must end). Volume weighting decides the cold-start constitution of Case 8.1: requiring many transactions before a score displays protects buyers from noise and walls newcomers out, so platforms invent probationary architectures - 'new seller' badges, escrow, platform-guaranteed first transactions - that substitute institutional trust while personal reputation accumulates. Display thresholds create regulatory cliffs: where a 4.6 is invisible and a 4.7 is 'top rated', the threshold is a licensing regime, and the anxiety it induces

in Chapter 7's workers is not incidental but the disciplinary function of Section 8.1 given a number. The designer's notebook closes where the chapter did: every lever - elicitation, reveal timing, dimensionality, decay, thresholds - redistributes power among buyers, sellers, incumbents, and entrants, and the platform that claims its ratings merely reflect quality is claiming, in the terms of Section 8.3, that commensuration is not a social process. Discussion questions: (1) Choose a platform and reverse-engineer its reputation constitution across the five levers; whom does each setting favour? (2) Design the reputation system for a market in Chapter 7's crowdwork, where the *requesters* are the party needing discipline; which levers move, and why? (3) Should reputation decay be a legal right - a commercial right to be forgotten? Argue from Section 8.3 and Chapter 10's data-protection architecture.

8.7. In Depth: From Persons to Systems - The Sociology of Institutional Trust

Sections 8.1 and 8.6 analysed trust between transacting strangers; this section ascends a level, to trust in the *systems* - platforms, protocols, professions, states - within which those transactions run, because the digital economy's deepest trust question is not whether the seller will ship but whether the infrastructure that scores the seller deserves the reliance placed on it.

Two productions of trust

The sociological baseline is Zucker's (1986) historical account of how American economic modernization produced trust when personal acquaintance could no longer carry exchange: *process-based* trust (reputation from past dealings), *characteristic-based* trust (similarity - the homophily that Section 4.4 met as a filter), and, decisively for scale, *institution-based* trust - trust manufactured by formal structures: professional credentials, regulation, insurance, escrow, the legal system as guarantor of last resort. Zucker's thesis was that nineteenth-century America, flooded with strangers and distance, substituted the third mode for the first two - and the digital economy, this section argues, ran the substitution again in compressed time, with a twist the theory illuminates. The platforms of Chapter 3 did not merely build process-based trust (the ratings of Section 8.1); they built the *institutions* - escrow, guarantees, identity verification, dispute resolution (Section 3.7) - privately, so that institution-based trust, historically the province of professions and states, became a proprietary product. The user who 'trusts a stranger's car' trusts, in fact, an insurance-and-rating apparatus; the confidence is real, and its object is a firm.

The reflexivity problem and the new intermediaries of doubt

The ascent has a second storey. Institution-based trust presupposes trust in the institutions, and here the digital economy exhibits its characteristic reflexive twist: the systems that manufacture interpersonal trust - ratings, feeds, scores, verification - are themselves opaque (Section 6.2), conflicted (Section 3.7), and gameable (Section 8.6), so the doubt migrates upward, from counterparties to infrastructures. The pattern is visible across the book's cases: review fraud converts trust in strangers into distrust of the rating system; misinformation converts distrust of claims into distrust of the informational order (Case 9.2's deepest finding). Each migration summons a new stratum of intermediaries of doubt - verification vendors, fact-checkers, audit firms, the algorithm auditors of Chapter 10 - whose own trustworthiness then becomes the question, in a regress that certification theory (Section 2.8) taught us to expect. Who audits the auditor is not a rhetorical flourish but the recursive structure of trust at scale.

Two institutional responses close the section and the chapter's arc. The first is the *trustless* ambition of cryptographic systems - replace institutions with verification, the ledger with mathematics - whose instructive career demonstrated the sociology it hoped to escape: the protocols minimized one trust relation and re-concentrated others (exchanges, wallets, developers, oracles), because, as this chapter has argued throughout, exchange among humans cannot be de-institutionalized, only re-institutionalized. The second is the public-law response of Chapter 10 - audit rights, transparency obligations, due-process mandates - which is, in Zucker's terms, the state re-entering the production of institution-based trust for infrastructures that had privatized it. The sociology of digital markets thus ends where the book's argument has been heading from Chapter 1: trust, like scarcity, like classification, like the market itself, is manufactured - and the question is never whether, but by whom, under what audit. Discussion questions: (1) Classify five trust devices from earlier chapters (escrow, verified badge, rating, insurance, blockchain settlement) into Zucker's three modes; which mode dominates the digital economy, and why? (2) Trace one full migration of doubt - from counterparty to system to auditor - in a market you know, and identify where the regress currently rests. (3) 'Trustlessness is re-institutionalization in denial.' Defend and attack with the cryptocurrency record.

Chapter 9. Media, Attention, and the Political Economy of Communication

Learning objectives. After studying this chapter, the reader should be able to: (1) define the attention economy and derive its dynamics from the scarcity theorem of Simon and the audience economics of the advertising model; (2) analyse the platformization of media industries, with journalism as the critical case, and evaluate the resulting dependencies between publishers and platforms; (3) present the theory of the public sphere and assess how algorithmic curation transforms it; (4) evaluate the empirical literature on filter bubbles, echo chambers, and the spread of misinformation, distinguishing documented mechanisms from popular exaggerations; and (5) explain why content moderation is a constitutive, not incidental, activity of platforms, and what follows for their claimed neutrality.

KEY CONCEPTS

attention scarcity; advertising markets; platformization; recommender systems; ranking; misinformation; metrics; referral traffic; generative search; synthetic media; provenance

This chapter completes the book's third lens. The economic analysis of Chapters 2 through 5 established that the dominant consumer platforms monetize attention; the present chapter asks what happens to media - the institutions that inform, entertain, and constitute publics - when attention becomes the master commodity and platforms become its market-makers. The stakes exceed industry economics: the media system is the infrastructure of democratic self-understanding, and its reorganization by platform logic is therefore a political fact of the first order.

9.1. The Attention Economy and Advertising Markets

The founding theorem was stated before the Internet existed. Simon (1971), observing the coming information abundance, drew the inference that most commentary still resists: what information consumes is the attention of its recipients, so a wealth of information creates a poverty of attention - scarcity does not disappear in the information economy; it migrates from content to audience. Goldhaber (1997) drew the economic corollary at the web's dawn: when information is abundant and reproducible at zero cost (Chapter 2), attention becomes the binding constraint and, increasingly, the currency - sought, measured, traded, and hoarded. Wu's (2016) history supplies the industrial genealogy: the 'attention merchants' - from the penny press through broadcasting to the feed - share a single business architecture, the resale

of assembled audiences to advertisers (Smythe, 1977), refined at each stage by better capture and better measurement, and perfected in the targeted, auctioned, individually attributed advertising of Chapter 5.

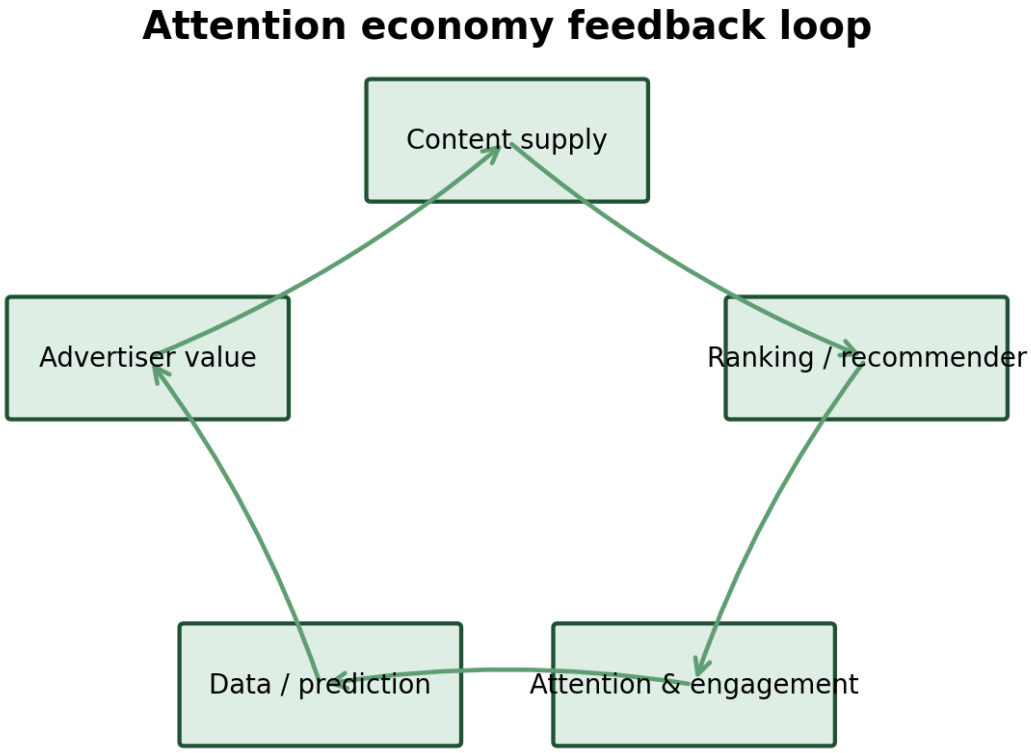

Figure 9.1. The attention-economy feedback loop linking ranking, engagement, data and advertiser value.

Webster (2014) supplies the missing microstructure: the *marketplace of attention* is jointly produced by audiences, whose finite time and habitual repertoires set the supply; by media, which compete to capture it; and - the decisive digital addition - by *measurement regimes*, the metrics and algorithms that make attention visible, comparable, and saleable, and which therefore do not merely record the competition but referee it. Three structural consequences organize the rest of the chapter. First, attention markets concentrate: audience formation exhibits the winner-take-most dynamics of Chapter 3, since popularity is itself informational and recommendation amplifies it. Second, competition for a fixed budget of attention is zero-sum in a way competition for money is not - every minute won is won from somewhere - which rationalizes the engagement engineering, the autoplay, the infinite scroll, and the notification economy as competitive necessities rather than moral failings of individual firms. Third, what maximizes engagement is not what maximizes accuracy, deliberation, or welfare - arousal, outrage, novelty, and confirmation perform reliably - so the attention economy contains a standing bias whose civic consequences Section 9.3 examines. The

economist will recognize an externality; the media scholar, a transformation of the public sphere; both are correct, and Chapter 10 asks what either can do about it.

COMMON MISCONCEPTION

More engagement is not the same as more persuasion, and more exposure is not the same as more influence. Research on misinformation must distinguish production, reach, selective exposure, memory, belief change, and behaviour rather than treating "viral" content as a direct measure of social effect.

9.2. Platformization of Media Industries and Journalism

Nieborg and Poell's (2018) concept of platformization, introduced for cultural production in Chapter 5, applies to media industries with particular force: news, music, video, and publishing have been reorganized around platform infrastructures for distribution, around platform metrics for editorial judgment, and around platform markets for monetization. Journalism is the critical case because its economics were a historical accident that digitalization dissolved. The twentieth-century newspaper was a bundle - civic reporting cross-subsidized by classified advertising, sport, and listings - held together by control of local distribution. The Internet unbundled it on every seam: classifieds departed to specialized marketplaces (Chapter 5), readers' attention departed to feeds and search, and the advertising that followed the attention was resold by platforms with better targeting (Chapter 5) and none of the newsroom's costs. Pickard (2020) states the outcome bluntly: the commercial model of journalism has collapsed for structural, not managerial, reasons - the market failure is genuine, the civic information it funded is a public good in the Chapter 2 sense, and the policy question is no longer whether but how non-market support should be organized.

The relationship between the survivors and the platforms is one of asymmetric dependence, and Nielsen and Ganter (2018) supply its canonical study: publishers cannot reach audiences without platform distribution, yet the terms of that distribution - the algorithm's weightings, the product formats, the revenue splits - are set unilaterally and changed without notice, so that editorial organizations invest in formats (video pivots, instant articles) that a single ranking change can strand. The structure is precisely the complementor's predicament of Chapter 3, with a civic multiplier: when the enveloping platform's priorities shift, what is stranded is not merely a product line but a community's information supply. Regulatory responses - bargaining codes obliging platforms to pay for news, public subsidies, nonprofit models - are examined in Chapter 10; here the analytical point is that they are all attempts to re-price, by institutional force, an externality the attention market does not see: journalism's value to non-readers, which is to say its democratic value.

9.3. Algorithmic Curation, Public Spheres, and Misinformation

The normative benchmark against which these transformations are measured descends from Habermas (1989): the public sphere as the social space in which private citizens reason together about common affairs, historically constituted by the very press whose economics Section 9.2 buried. The platform-era question is whether algorithmically curated feeds can host such a sphere, and Habermas's (2022) own late reflection states the sceptical case: platforms dissolve the editorial *gatekeeping* that, for all its exclusions, guaranteed a common agenda and professional filtering; in its place they install an attention-optimizing selection that fragments publics into self-referential spaces and blurs the line between public reasoning and private expression. The theory found its popular vocabulary in Pariser's (2011) *filter bubble* - personalization, invisible and unchosen, would enclose each citizen in a flattering informational cocoon.

The empirical literature obliges the textbook writer to complicate the popular story. Bakshy et al. (2015), with platform data, found that algorithmic ranking does reduce exposure to cross-cutting content - but by less than users' own choices of friends and clicks do; Flaxman et al. (2016) found that social and search channels are associated with *higher* ideological segregation of news consumption than direct browsing, and simultaneously with *more* exposure to opposing perspectives - polarization and diversity rising together, against the bubble metaphor's prediction. The honest synthesis is that echo chambers are real but minoritarian, self-selection dominates algorithmic selection, and the bubble is better understood as a feedback between the two than as a technological cage. Misinformation research imposes the same discipline. Vosoughi et al. (2018), analysing a decade of rumour cascades, found false news travelling significantly farther, faster, and deeper than truth - propelled by novelty and emotional arousal, and by humans more than by bots: the attention economy's selection bias (Section 9.1), operating through people. Wardle and Derakhshan (2017) supply the field's taxonomy - misinformation, disinformation, and malinformation, distinguished by falsity and intent - and its ecological insight: the problem is an *information disorder* of the whole system, not an inventory of false items to be deleted.

Large-scale experiments conducted with platform cooperation have since tested the algorithmic hypothesis directly. During the 2020 United States election campaign, Guess et al. (2023) moved a random sample of users from algorithmically ranked to reverse-chronological feeds for three months. Exposure changed substantially - users saw more political and more untrustworthy content and spent less time on the platforms - yet polarization and political

attitudes did not change detectably. Deactivation experiments point in a complementary direction. Paying users to stop using a social network for four weeks reduced their news knowledge but also their political polarization, and it increased self-reported well-being (Allcott et al., 2020). The staggered introduction of a social network across American college campuses was followed by a deterioration in students' mental health (Braghieri et al., 2022). Taken together, the evidence suggests that platforms affect their users less through the ranking formula alone than through the time, attention and social comparison they absorb, which shifts the policy debate of Chapter 10 from the feed's algorithm toward the design of the service as a whole.

Which returns the chapter, and the book's media lens, to moderation. Gillespie (2018) states the constitutive thesis that Chapter 7's account of moderation labour (Roberts, 2019) prepared: moderation is not an unfortunate addendum to platforms but the essence of what they offer - every feed is a selection, every selection embodies values, and the platforms are therefore *custodians* of public discourse whether they accept the title or not. Napoli (2019) draws the regulatory inference: institutions that perform the gatekeeping functions of media while disclaiming the responsibilities of media occupy a category that communication policy has always, eventually, brought within public-interest obligations - as it did broadcasting - and the platform era's defining media-policy question is what those obligations should be when the gatekeeper is an auction-driven algorithm. That question is continuous with the competition, data, and speech regulation of Chapter 10, where this book's three lenses finally converge on the same object: the platform as market, as institution, and as medium.

WORKED EXAMPLE

Scenario. A publisher receives 10 million monthly page views. Search provides 40% of referrals. A change in search interfaces reduces search referrals by 25%, with no immediate replacement channel.

Calculation. Search-originated views fall from 4 million to 3 million, reducing total monthly page views to about 9 million: a 10% total traffic decline.

Interpretation. The example shows how an apparently external interface change can alter newsroom economics even when the publisher's own product, audience preferences, and editorial output are unchanged.

THREE LENSES

Economics asks: How is scarce attention priced, auctioned, bundled, and converted into revenue?

Sociology asks: How do networks, identities, and social position shape exposure, trust, and participation?

Media studies asks: How do ranking, recommendation, generative interfaces, and metrics restructure gatekeeping and public visibility?

2026 UPDATE: GENERATIVE SEARCH AND THE NEW GATEWAY PROBLEM

The platformization of news now includes generative interfaces that can answer a user's question before a click reaches the publisher. The Reuters Institute's Digital News Report 2026 identifies declining search referrals as a central strategic concern and documents the emerging use of AI chatbots as a route to news (Reuters Institute for the Study of Journalism, 2026).

This creates a new form of intermediation: the platform is not only ranking links but synthesizing responses. The economic issues are referral substitution, licensing, attribution, bargaining power, and conversion; the media issues are provenance, source visibility, error correction, and the weakening or redesign of the publisher-audience relationship.

For analysis, distinguish recommendation systems (which rank available items) from generative answer systems (which can transform, summarize, and recombine source material into a new interface object).

9.4. Chapter Summary and Discussion Questions

This chapter analysed the digital economy as a media order. The attention economy was derived from Simon's (1971) scarcity theorem and Goldhaber's (1997) corollary, given industrial history by Wu (2016) and microstructure by Webster (2014), and shown to carry a structural bias toward engagement over deliberation. Platformization was applied to journalism: the dissolution of the advertising bundle, the genuine market failure of civic information (Pickard, 2020), and the asymmetric publisher-platform dependence (Nielsen & Ganter, 2018) analysed as the complementor problem with democratic stakes. The public-sphere tradition (Habermas, 1989, 2022) framed algorithmic curation, and the empirical record was presented with its complications: modest algorithmic effects relative to self-selection (Bakshy et al., 2015;

Flaxman et al., 2016), the novelty-driven advantage of falsehood (Vosoughi et al., 2018), and the systemic framing of information disorder (Wardle & Derakhshan, 2017). Moderation was established as constitutive (Gillespie, 2018) and the public-interest question posed (Napoli, 2019).

Discussion questions. (1) Restate Simon's theorem and use it to explain one concrete design feature of a feed you use. What would the feature look like if the service were subscriber-funded (Chapter 5)? (2) Was the newspaper bundle a market solution or a historical accident? What follows from each answer for journalism policy? (3) Using Nielsen and Ganter (2018), write the risk section of a publisher's strategy memo on platform dependence. Which risks can be hedged, and which are structural? (4) The filter-bubble thesis and the Flaxman et al. (2016) findings appear to conflict. Reconcile them, specifying the mechanisms each captures. (5) Vosoughi et al. (2018) found falsehood outruns truth among humans, not only bots. What interventions does this finding favour, and which popular remedies does it undercut? (6) 'Platforms are media companies.' Prosecute and defend, using Gillespie (2018) and the platforms' own terms of service as evidence, and state what regulatory consequences hang on the verdict.

9.4a. Numerical and Data Exercises

Exercise 9.A (a keyword auction). Three advertisers bid EUR 3.00, EUR 2.00 and EUR 1.00 per click for two ad positions with click-through rates of 10% and 5%. Quality scores are equal. (a) Under a generalized second-price auction, who wins which position and what does each pay per click? (b) What is the platform's expected revenue per 1,000 impressions? (c) How would the outcome change in principle if the lowest bidder raised its quality score enough to rank second (Section 5.7)?

Exercise 9.B (advertising versus subscription). A news site has 1.2 million monthly users who view 4 pages each on average; display advertising earns EUR 5 per 1,000 page views. A metered paywall converts 2.5% of monthly users into subscribers paying EUR 7 per month but reduces page views by 20%. (a) Compute monthly advertising revenue before and after the paywall. (b) Compute monthly subscription revenue. (c) What does the result imply for the incentives described in Section 5.1 and Case 5.2?

Indicative answers. 9.A (a) The EUR 3.00 bidder takes the first position and pays EUR 2.00 per click; the EUR 2.00 bidder takes the second and pays EUR 1.00. (b) 1,000 × (0.10 × 2.00 + 0.05 × 1.00) = EUR 250. (c) The quality-weighted ranking would give the former lowest bidder

the second position at a price set by the rival's quality-weighted bid; the platform's quality scores thus co-determine both allocation and price. 9.B (a) Before: 4.8 million views, EUR 24,000; after: 3.84 million views, EUR 19,200. (b) 30,000 subscribers × EUR 7 = EUR 210,000. (c) Where readers are willing to pay, subscription revenue dwarfs display advertising, so the outlet's incentives shift from traffic toward retention and perceived value.

9.5. Case Studies and Debates

The materials below carry the chapter into a state-versus-platform confrontation over the price of news, the election that made misinformation a research field, and the liability question on which platform regulation ultimately turns.

Case 9.1. The week the news disappeared: Australia's bargaining code

Section 9.2 described publisher-platform dependence as asymmetric; Australia's News Media Bargaining Code, enacted in 2021, was the first attempt by a state to reverse the asymmetry by statute - and the confrontation it provoked measured the asymmetry publicly. The code's design responded precisely to the bargaining problem of Nielsen and Ganter (2018): designated platforms would be required to negotiate payment with news businesses for the value of news in their services, with final-offer arbitration as the backstop that made refusal to deal costly. The platforms' responses diverged instructively. One negotiated. The other - the dominant social network - demonstrated its outside option: for several days it removed news for Australian users entirely, a blackout whose blunt implementation also swept away pages of health agencies and emergency services, before restoring news in exchange for amendments softening the designation mechanism. Deals followed on both platforms - tens of millions of dollars annually flowing to Australian publishers, large and small - and the code became the template studied and imitated, with variations, from Canada to California. Its critics, meanwhile, assembled the counter-ledger: payments negotiated in private, skewed toward the largest incumbent publishers, terminable at the platforms' discretion - as later non-renewals demonstrated - and premised on a contested theory of who confers value on whom (does the feed take the news's value, or give the news its reach?).

Seminar questions: (1) Model the pre-code bargaining game between a publisher and a platform, and identify exactly which parameter final-offer arbitration changes. (2) The blackout showed the platform could exit news at low cost; what does that imply about the direction of value flow - and does the answer differ for the search engine? (3) Compare three instruments for funding journalism - bargaining codes, taxes on platform advertising earmarked for news, and direct public subsidy - against Pickard's (2020) market-failure diagnosis and against capture risk. (4) Payments went disproportionately to large publishers; design an allocation mechanism that reaches the local journalism whose loss Section 9.2 identified as the civic core of the problem.

Case 9.2. 2016 and the making of a research field

The misinformation literature of Section 9.3 has a founding shock: the 2016 United States presidential campaign, during which fabricated news stories - manufactured by ideological entrepreneurs and, notoriously, by profit-seeking teenagers abroad optimizing for the advertising revenue of Chapter 5 - achieved massive circulation through exactly the engagement machinery of Section 9.1. Allcott and Gentzkow (2017) supplied the field's disciplining first estimate. Compiling databases of stories verified as false, they measured circulation and exposure: pro-Trump fabrications were shared far more than pro-Clinton ones; yet survey-based exposure estimates suggested the average American adult saw on the order of one to several fake stories, with recall and belief lower still, and social media, while the fabrications' principal habitat, remained a minority source of election news overall. Their back-of-envelope persuasion calculation - for fake news to have changed the outcome, a fake story would have needed a persuasive effect implausibly larger than that measured for conventional campaign advertising - became the canonical caution against the strongest causal narrative.

The case teaches method as much as substance, and pairs deliberately with Section 9.3's other findings: Vosoughi et al. (2018) established that falsehood *travels* farther and faster; Allcott and Gentzkow established that travelling is not persuading; and the field matured in the space between - toward supply-side studies of the fabrication economy, demand-side studies of why sharing is expressive rather than epistemic, and systemic framings (Wardle & Derakhshan, 2017) in which the deepest damage is not false belief implanted but true institutions distrusted: the 'fake news' epithet itself was captured, within a year, as a weapon against journalism. Seminar questions: (1) Reconstruct the Allcott-Gentzkow persuasion arithmetic and attack its weakest assumption. (2) Reconcile 'falsehood diffuses faster' with 'exposure and persuasion were modest' - which quantity does each study measure, and which matters for which policy? (3) The fabricators' motive was advertising arbitrage; trace the value chain from Section 5.1 and identify the intervention point - ad networks, ranking, or user friction - with the best cost-benefit profile. (4) If misinformation's principal harm is corrosion of trust rather than false belief, which of Chapter 10's instruments still bite?

Debate 9.1. Amplification and liability: who answers for the feed?

Motion: platforms should bear legal liability for content their recommendation systems algorithmically amplify, even where they bear none for merely hosting it. *For the motion:* the intermediary-immunity settlements of the 1990s - shields protecting hosts from liability for users' speech - answered a hosting question in a hosting era. The feed changed the object, for ranking is the platform's own conduct, not the user's: the aggregator of Section 5.2 selects,

orders, and promotes for engagement, exercising precisely the editorial function (Gillespie, 2018) whose human analogue - the publisher - has always answered for foreseeable harms. Distinguishing hosting from amplification therefore preserves the open Internet's protection for speech while attaching responsibility to the machinery that monetizes its virality, and the systemic-risk obligations of the Digital Services Act (Chapter 10) already regulate the recommender as conduct in exactly this spirit. *Against the motion:* the hosting/amplification line dissolves under litigation pressure - every ordering of infinite content is 'amplification' of something, so liability for ranked harm is liability for ranking as such, and the rational platform response is over-removal at the scale where false positives are millions of lawful posts - a privatized censorship falling, per Section 8.2's gradients, heaviest on marginal speakers. Causation is unprovable in the individual case (which recommendation caused which act?), inviting liability standards that are really licensing regimes. And the empirical foundation - that algorithmic amplification, rather than demand-side selection (Section 9.3), drives the harms - is precisely what the careful literature failed to confirm. The debate should require each side to draft its operative legal test in one sentence, apply it to three items - a defamatory post ranked highly, a terrorist recruitment video recommended, and an eating-disorder spiral of lawful content - and accept, on the record, the errors its test produces.

9.6. In Depth: Metrics at Work - How Measurement Remade the Newsroom

Section 9.1 identified measurement regimes as the referees of the attention marketplace; this advanced section follows the metrics inside the institution this chapter cares most about, because journalism is where the reactivity of Section 8.3 - the measured adapting to the measure - collides most directly with a public function.

From the ratings book to the real-time dashboard

Audience measurement has always constituted the audience it claimed to describe: broadcast ratings and print circulation audits were the currencies in which Section 9.1's attention merchants denominated their inventory, and their methodological choices - who counts as a viewer, which diaries and meters, which demographics command premiums - shaped programming for decades (Webster, 2014). What digitalization changed was granularity, latency, and address: the newsroom that once learned its audience quarterly, in aggregate, now watches a real-time dashboard that scores each article, each headline variant, each minute of attention - and the metric arrives not as an external audit but as an ambient presence on screens above the desks. The stakes follow from the chapter's economics: under the advertising model of Section 5.1, the dashboard's currency is the business model, so the question of how journalists metabolize metrics is the question of how the attention economy's selection bias (Section 9.1) enters editorial judgment itself.

The ethnographic finding: same numbers, different newsrooms

The essential study is Christin's (2020) comparative ethnography of web newsrooms in the United States and France - organizations using the same analytics tools whose journalists nonetheless lived with them differently. Christin's central finding disciplines both technological determinism and professional complacency: metrics are refracted through occupational cultures and organizational buffers. In the American newsroom, traffic was morally loaded but publicly disavowed - journalists performed indifference to clicks while compensation, commissioning, and anxiety tracked them - a regime she calls one of *disciplined* engagement, where the dashboard governs through internalization. In the French newsroom, an older professional hierarchy held the metric at arm's length - traffic belonged to a subordinate caste of web producers while prestige remained with print-descended criticism - buffering editorial judgment at the cost of strategic drift. Neither newsroom escaped the number; each institutionalized it differently, and the difference mattered for what was published. The generalization is the section's payload: the reactivity of metrics (Section 8.3) is not a mechanical effect but an

organizational accomplishment, which means it is a design variable - newsrooms, like the platforms of Section 8.6, choose their measurement constitutions, deciding which metrics are visible to whom, attached to what rewards, and buffered by which professional norms.

The synthesis returns to the chapter's civic register. Subscription-funded outlets (Case 5.2) rebuilt their dashboards around retention and subscriber conversion rather than raw traffic - and editorial behaviour followed the new currency, toward depth, habit, and the loyal reader, vindicating Section 5.1's claim that the revenue model selects the metric and the metric selects the journalism. The public-interest question of Section 9.3 thereby acquires an institutional answer alongside the regulatory one: if what is measured disciplines what is made, then the governance of news metrics - who defines engagement, whether civic reach counts, what the dashboard renders invisible - is media policy conducted by interface design. Discussion questions: (1) Use Christin (2020) to argue against the claim that 'metrics inevitably produce clickbait'; what, exactly, mediates? (2) Design the analytics dashboard for a public-service newsroom: which three metrics appear, which are deliberately hidden, and what does each choice cost? (3) Connect the visibility thresholds of Section 8.6 to a recommendation feed's treatment of news outlets; where is the 'top rated' cliff in journalism, and who sets it?

9.7. In Depth: The Pipes - A Political Economy of Internet Infrastructure

This chapter has analysed the attention layer; this section descends the stack, because the feeds ride on physical and quasi-physical infrastructures - cables, exchanges, content-delivery networks, interconnection agreements - whose economics and politics condition everything above them, and whose governance DeNardis (2014) taught the field to read as politics conducted by technical means.

The stack below the platforms

The Internet's transport economy has three strata the media scholar should know as well as the economist. At the bottom, the *physical* layer: submarine cables carry the overwhelming share of intercontinental traffic, their routes retracing - with instructive fidelity - the telegraph geography of an earlier globalization, their ownership migrating from telecommunications consortia to the platform firms themselves, which now finance private cables to join their own data centres: vertical integration reaching literally to the seabed, with the sovereignty implications of Section 10.3 attached (cables are cut, tapped, and landing-licensed by states). In the middle, *interconnection*: the Internet is a network of networks stitched together by peering (settlement-free exchange between comparable networks) and transit (paid carriage), governed not by regulation but by a thicket of private agreements whose disputes - who pays whom when traffic is asymmetric - periodically surface as degraded service and political conflict, the recurring 'fair share' campaigns by carriers to charge content providers being the current instalment of a very old argument about who funds the pipes. At the top of the transport economy, the *content-delivery* layer: caches and delivery networks that position popular content near users, converting Chapter 2's zero transportation cost from theory into engineering - and concentrating, in a handful of CDN and cloud providers, a chokepoint whose outages take large fractions of the visible web offline at once and whose deplatforming decisions constitute, in effect, infrastructure-layer moderation: the custodianship of Gillespie (2018) exercised below the application, with fewer procedures and less visibility.

Net neutrality and the politics of the middle

The stack's most durable policy war concerns the access networks' treatment of the traffic they carry. Wu (2003) named the principle: *network neutrality* - the access provider should not discriminate among applications and sources - grounding it in the innovation economics this book has assembled. The end-to-end architecture of Section 1.11 made the network a neutral platform on whose edges permissionless innovation flourished (the external

enabler of Chapter 4), and discrimination by carriers - blocking, throttling, paid prioritization - would convert the transport layer into a gatekeeper collecting the rents of Section 3.4 while taxing exactly the entrants that contestability requires. The opposing economics is two-sided-market theory turned against its teachers: carriers argue that charging content providers is the efficient price structure of Section 3.2 (the subsidy side re-priced), that prioritization is quality differentiation serving latency-sensitive uses, and that neutrality mandates suppress the network investment on which everything rides. The regulatory history - rules adopted, repealed, readopted, and litigated across jurisdictions and decades, with the European framework comparatively stable and the American one oscillating with administrations - is itself the lesson: the question is not settled because it is genuinely the book's central question in miniature, *who may tax the flow between producers and publics, under what accountability* - asked, this time, of the pipes rather than the feeds. The section closes the media lens's arc: the political economy of communication runs, uninterrupted, from the submarine cable to the ranking algorithm, and the student who has descended the stack once will not again mistake the feed for the whole of the medium. Discussion questions: (1) Map the three strata onto Chapter 3's concepts: where are the network effects, the switching costs, the envelopment opportunities? (2) Reconstruct the 'fair share' dispute as a two-sided pricing problem; who is the subsidy side, and what does Wu's (2003) innovation argument add that pricing theory omits? (3) A CDN terminates service to a lawful but odious website; compare this act with a platform's moderation decision along visibility, procedure, and remedy - and say which layer's custodianship worries you more, and why.

Chapter 10. Regulation, Governance, and the Future of the Digital Economy

Learning objectives. After studying this chapter, the reader should be able to: (1) explain why digital markets strain the consumer-welfare framework of modern competition policy and present the reform positions in the ensuing debate; (2) describe the European regulatory architecture - data protection, gatekeeper obligations, and platform accountability - and the mechanism by which it acquires global reach; (3) analyse Internet and platform governance as a contest among regulatory models and among states, firms, and civil society; (4) situate the digital economy in global development, including the position of countries outside the platform duopoly of power; and (5) construct disciplined scenarios for the coming decade using the book's analytical apparatus.

KEY CONCEPTS

competition policy; gatekeeper; interoperability; data protection; DMA; DSA; AI Act; NIS2; sovereignty; regulatory capacity; compliance; enforcement

Every chapter of this book has ended at a door marked regulation: the property settlements of Chapter 2, the concentration of Chapter 3, the classification struggles of Chapters 6 and 7, the discrimination and stratification of Chapter 8, the public-interest question of Chapter 9. This final chapter walks through the door. Its premise is Lessig's (1999): behaviour in digital environments is regulated by four modalities - law, norms, markets, and *code*, the architecture that makes some conduct impossible and other conduct default - and the platform era is best understood as the discovery that whoever writes the code performs, without election, the work of the other three. Public regulation of the digital economy is therefore not the intrusion of politics into a natural market; it is the renegotiation of a division of regulatory labour that private actors have already, and profitably, assumed.

10.1. Competition Policy for Digital Markets

Modern antitrust enforcement, particularly in the United States, was rebuilt in the late twentieth century around the consumer-welfare standard associated with Bork (1978): the law's sole legitimate concern is efficiency as reflected, principally, in consumer prices, and concentration achieved through superior product is not an offence but the reward of competition. The digital economy stress-tests every load-bearing element of this framework, for reasons the reader now possesses. Where the diagnostic instrument is price, zero-price two-

sided markets (Chapter 3) render it silent: the consumer pays nothing, pays in attention and data, and may be harmed in quality, privacy, and choice while the price gauge reads zero. Where the presumption is that monopoly invites entry, network effects, switching costs, and data feedback loops (Chapters 3 and 6) make incumbency self-reinforcing. And where merger review asks whether today's overlap raises today's prices, the characteristic digital acquisition removes tomorrow's rival - the nascent competitor bought before it becomes one - a harm the framework can name only with difficulty.

Khan's (2017) article on the antitrust paradox of the platform era crystallized the reform position now labelled Neo-Brandeisian: predatory below-cost expansion and vertical envelopment (Chapters 3 and 5) are rational, durable strategies for platforms in a way Chicago-school theory said they could not be. The marketplace that competes with its own complementors (Zhu & Liu, 2018) presents conflicts of interest that price analysis cannot see. Antitrust, on this view, should recover its older concerns with structure, dependence, and private power - Wu (2018) supplies the historical argument that it once had them. The institutional response has been broadly convergent across expert reviews: the Crémer report for the European Commission (Crémer et al., 2019) and the Furman review for the United Kingdom (Furman et al., 2019) each concluded that case-by-case competition law is too slow for tipping-prone markets and recommended ex ante regimes - standing conduct rules for firms with entrenched gateway positions, error costs recalibrated on the premise that under-enforcement, not over-enforcement, is the digital default. Whether this constitutes the modernization of competition policy or its transformation into sectoral regulation is the live doctrinal question; that the pure consumer-welfare toolkit was insufficient for the market structures this book has described is, by the 2020s, close to an official consensus outside the United States and a contested one within it.

The empirical case for concern about acquisitions has also strengthened. In the pharmaceutical industry, where project-level data make the practice observable, Cunningham et al. (2021) show that incumbents sometimes acquire innovative targets in order to discontinue overlapping projects - 'killer acquisitions' that account for a non-trivial share of deals. In digital markets, where acquired products are more often integrated than discontinued, Kamepalli et al. (2020) model a subtler 'kill zone': when users expect a dominant platform to buy or copy successful entrants, early adopters hesitate to switch to them, and venture investors reduce funding for startups in the platform's adjacent space. Both mechanisms supply a theory of harm that the consumer-welfare framework found hard to state, and both inform the DMA's

requirement that gatekeepers inform the Commission of all intended acquisitions in the digital sector.

COMMON MISCONCEPTION

Digital regulation is not a single trade-off between "innovation" and "control." Different instruments target different market failures and rights: competition rules address market power, data protection addresses lawful processing and rights, platform rules address systemic intermediary duties, cybersecurity rules address resilience, and AI rules target risks tied to AI systems and models.

10.2. Data Protection, Platform Regulation, and the European Model

The ex ante turn found its fullest expression in the European Union, whose regulatory architecture now stands as the world's most developed - and most exported - model. Its foundation is the General Data Protection Regulation (Regulation (EU) 2016/679), which constitutionalizes the data economy of Chapter 6: processing personal data requires a lawful basis; data subjects hold rights of access, erasure, and portability - the last aimed directly at the switching costs of Chapter 3; obligations follow the data across borders; and sanctions scale with global turnover. Upon this foundation the Union built the twin regulations of 2022. The Digital Markets Act (Regulation (EU) 2022/1925) writes the Crémer-Furman logic into law: designated *gatekeepers* - platforms controlling durable gateway positions - face standing obligations and prohibitions that read as a checklist of this book's Chapter 3 and 5 mechanisms: no self-preferencing in ranking, no combining data across services without consent, interoperability and portability mandates, and freedom for business users from anti-steering restraints. The Digital Services Act (Regulation (EU) 2022/2065) addresses Chapter 9's questions: graduated due-diligence duties for intermediaries, transparency of recommender systems and advertising, due process in moderation, and, for the largest platforms, obligations to assess and mitigate *systemic risks* to civic discourse - the public-interest obligations Napoli (2019) anticipated, drafted in the idiom of risk regulation.

A further layer governs data itself rather than the firms that hold it. The Data Governance Act creates a framework for the re-use of protected public-sector data, for neutral data-intermediation services and for 'data altruism' (European Parliament & Council of the European Union, 2022c). The Data Act, applicable since September 2025, gives users of connected products - from cars to industrial machinery - a right to access the data their use generates and to share it with third parties; it also restricts unfair contractual terms on data sharing imposed on smaller firms and obliges cloud providers to facilitate switching (European Parliament & Council of the European Union, 2023b). In the terms of this book, the Data Act

turns the Jones-Tonetti argument of Section 6.7 into law: rather than leaving non-rival data hoarded by manufacturers, it reallocates access rights to users so that the data can serve more uses. Whether it achieves that aim or mainly adds compliance costs is among the empirical questions of the coming decade. The independent assessment of the DMA's design by Cabral et al. (2021) offers a template for such evaluations: it endorsed the architecture but warned that the information asymmetry between regulators and the regulated would remain the main obstacle to enforcement.

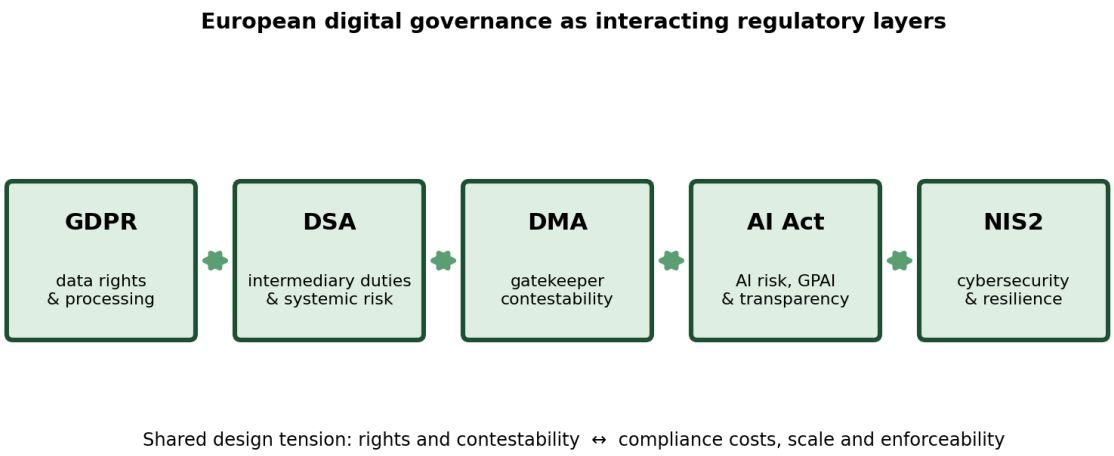

Figure 10.1. The interacting layers of European digital governance.

The model's global significance runs through what Bradford (2020) named the *Brussels effect*: firms serving the large, wealthy, regulation-capable European market standardize on its rules globally - compliance is indivisible or divisibility is costlier than uniformity - so that the Union exports regulation without treaty or coercion, as the spread of GDPR-style statutes across dozens of jurisdictions attests. Bradford (2023) generalizes the observation into the decade's geopolitical frame: three *digital empires* project competing governance models - an American market-driven model that trusts private ordering, a Chinese state-driven model that subordinates platforms to political authority, and the European rights-driven model - and most of the world's states, firms, and users live in the overlap, choosing and mixing under pressure. The framework disciplines two lazy narratives at once: regulation is not the opposite of the digital economy but one of its constitutive settlements (Chapter 2 said as much about property); and the settlement is plural - there is no single 'digital future', only futures under different constitutions of code.

10.3. Global Governance, Digital Sovereignty, and Development

Beneath state regulation lies the question of who governs the network itself. DeNardis (2014) makes the field's central point: Internet governance is exercised less in declarations than in *infrastructure* - protocols, standards bodies, the naming and numbering systems, interconnection agreements - where technical design decisions allocate rights and power without ever announcing themselves as politics. The multistakeholder institutions that steward these layers embody one governance model; the recurring pushes to relocate authority to intergovernmental bodies embody another; and the vocabulary of *digital sovereignty* - data localization, national clouds, filtered networks, indigenous platform champions - names the strategies by which states reassert territorial authority over an architecture designed to ignore it. The sovereignty turn is not one thing: it spans authoritarian information control, European strategic autonomy, and developmental protectionism, and the analyst must ask in each case *whose* autonomy is being secured, the state's or the citizen's.

Central and Eastern Europe offers instructive episodes of this politics. Hungary introduced a steeply progressive tax on advertising revenue in 2014. The European Commission found its progressive structure to be unlawful state aid, but the Court of Justice ultimately upheld the member state's freedom to design progressive turnover taxes (European Commission v Hungary, 2021). Poland's 2021 proposal for a levy on advertising revenue, presented partly as a way to tax global platforms, met a different kind of resistance: on 10 February 2021 nearly sixty independent media outlets suspended their normal output in a one-day protest, 'Media without choice', arguing that the levy would weaken domestic media already squeezed by those same platforms, and the proposal was not enacted. The two episodes show why digital taxation in smaller economies is simultaneously fiscal, competition and media policy - the three lenses of this book converging on a single instrument.

For the majority of the world's population, the governing fact is different: the digital economy arrives as infrastructure, platforms, and rules made elsewhere. UNCTAD (2019) documents the geography this book's mechanisms predict: the platform economy's market value is concentrated overwhelmingly in the United States and China; the data of the global South flows outward as raw resource and returns as services and intelligence priced abroad - the value-capture asymmetry that gives the data-colonialism thesis of Chapter 6 (Couldry & Mejias, 2019) its purchase. Yet digitalization simultaneously offers late developers genuine capabilities - the leapfrogging of mobile money, the market access of Chapter 4's born-global ventures and Chapter 7's remote work. The policy problem is therefore neither embrace nor refusal but the construction of terms: competition and data regimes with local teeth, investment in the skills of Chapter 8's divides, and South-South cooperation on infrastructure.

Development economics thus rejoins the book where Chapter 1 began: the definitional and measurement politics of the digital economy are, for most states, the politics of what they are allowed to see, tax, and negotiate.

10.4. Scenarios for the Next Decade

A textbook should not prophesy, but it can discipline speculation, and the book's apparatus yields three scenario axes. The first is *concentration versus contestability*: do the mechanisms of Chapters 3 and 6 - network effects, switching costs, data loops, now compounded by the scale economics of frontier AI - entrench the incumbent constellation, or do interoperability mandates, portability, open models, and the periodic platform shift (each wave of Chapter 1 unseated someone) reopen the field? The second is *the constitution of data and AI*: which of Bradford's (2023) three models - market-driven, state-driven, rights-driven - sets the effective global defaults for the technologies of Chapter 6, and does the Brussels effect extend to AI regulation as it did to data protection? The third is *the distribution of the surplus*: Acemoglu and Johnson (2023) supply the millennium-scale reminder that technological progress raises general living standards only when institutions - countervailing labour power, competition, the direction of innovation itself - steer it there, and their warning names the stakes of Chapters 6 and 7 precisely: an AI trajectory of 'so-so automation' plus surveillance is a choice, not a destiny, and so is its alternative, augmentation plus shared productivity. The scenarios multiply across the axes, but the method is the point: the reader who has followed this book holds the mechanisms - cost structures, network dynamics, institutional settlements, attention markets - from which any credible future of the digital economy must be assembled, and holds, too, the book's final lesson, which is that the assembly is performed not by technology but by the choices of firms, states, workers, and publics.

WORKED EXAMPLE

Scenario. A digital firm operates a marketplace, trains a general-purpose AI model, and provides a cloud service to regulated-sector clients.

Compliance map. The first analytical step is not to ask "Which digital law applies?" but to map activities: gatekeeping and self-preferencing; personal-data processing; systemic platform duties; GPAI obligations; cybersecurity and incident management; sector-specific client requirements.

Interpretation. Regulatory strategy is an architecture problem. Overlapping obligations attach to functions, roles, thresholds, and risks rather than to a company's label as a "tech firm."

THREE LENSES

Economics asks: Which market failure, externality, or rent does the intervention target, and what behavioural response should follow?

Sociology asks: Which institutions gain authority, whose rights become enforceable, and how is legitimacy produced?

Media studies asks: How do rules change moderation, ranking, access, transparency, and the governance of public communication?

2026 UPDATE: FROM REGULATORY DESIGN TO ENFORCEMENT

By 2026, the European digital rulebook has entered a materially different phase: the analytical question is no longer only what the DMA and AI Act require, but how obligations are interpreted and enforced.

On 23 July 2026, the European Commission issued two DMA non-compliance decisions against Google concerning self-preferencing in Search and anti-steering restrictions in Google Play, imposing fines totalling EUR 890 million (European Commission, 2026a).

From 2 August 2026, the AI Office and national authorities began exercising enforcement powers under the AI Act for provisions then applicable, including GPAI-related enforcement and transparency rules; the main high-risk-system obligations were deferred by the AI Omnibus (Regulation (EU) 2026/1744) to 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems (European Commission, 2026b; European Parliament & Council of the European Union, 2026).

Cyber resilience is also part of the governance architecture: NIS2 extends harmonized risk-management and reporting duties across a wider set of critical and digital sectors, linking market governance to operational resilience (European Commission, 2026c).

10.5. Chapter Summary and Discussion Questions

This chapter examined the governance of everything the book has described. Competition policy was shown to strain under zero prices, self-reinforcing incumbency, and nascent-competitor acquisitions, with the Neo-Brandeisian critique (Khan, 2017; Wu, 2018) and the expert-review consensus on ex ante regimes (Crémer et al., 2019; Furman et al., 2019)

presented against the consumer-welfare baseline (Bork, 1978). The European architecture - GDPR, Digital Markets Act, Digital Services Act - was mapped onto the book's mechanisms, and its global diffusion explained through the Brussels effect and the three-empires frame (Bradford, 2020, 2023). Internet governance was analysed as infrastructure politics (DeNardis, 2014) and the sovereignty turn disaggregated; the development question was posed through UNCTAD's (2019) geography of value capture. Finally, three scenario axes - contestability, the constitution of data and AI, and the distribution of the surplus (Acemoglu & Johnson, 2023) - organized the future as a set of institutional choices, completing the book's argument that the digital economy is, at every layer, a constructed and therefore contestable order.

Discussion questions. (1) Prosecute a zero-price platform under the consumer-welfare standard, then under a structural standard. Which harms can each name, and which must each ignore? (2) Take three Digital Markets Act obligations and trace each to the specific mechanism in Chapters 3, 5, or 6 it targets. Predict one evasion strategy per obligation. (3) Why does the Brussels effect operate for data protection but arguably not for labour law? Use the divisibility of compliance as your analytical variable. (4) 'Digital sovereignty' is claimed by authoritarian censors and European regulators alike. Construct criteria that distinguish the claims, and apply them to a current case. (5) A middle-income country asks you to draft its digital-economy strategy in five bullet points. Draft them, citing the relevant chapter mechanism for each, and identify the binding constraint you cannot legislate away. (6) Using Acemoglu and Johnson (2023), specify two institutional levers that would tilt AI toward augmentation rather than displacement, and the political economy that obstructs each.

10.5a. Numerical and Data Exercises

Exercise 10.A (the scale of sanctions). Under the DMA, fines may reach 10% of a gatekeeper's total worldwide turnover (20% for repeated infringements); under the GDPR, the maximum is 4%. For a firm with worldwide turnover of EUR 300 billion: (a) compute the maximum DMA fine for a first and for a repeated infringement, and the maximum GDPR fine; (b) express a EUR 890 million fine (Case 10.2) as a share of this turnover; (c) what determines whether a fine of this size deters?

Exercise 10.B (concentration). A market has four firms with shares of 60%, 25%, 10% and 5%. (a) Compute the Herfindahl-Hirschman Index (HHI). (b) Compute the HHI after a merger of the two smallest firms and the change in HHI. (c) Why might the HHI understate the competitive significance of a dominant platform acquiring a start-up with a 1% share (Section 10.1)?

Indicative answers. 10.A (a) EUR 30 billion; EUR 60 billion; EUR 12 billion. (b) 0.89 ÷ 300 ≈ 0.3% of turnover. (c) Deterrence depends on the expected fine - its size multiplied by the probability of detection and enforcement - relative to the profits from the conduct, and on remedies that change behaviour rather than only price it. 10.B (a) 3,600 + 625 + 100 + 25 = 4,350. (b) 3,600 + 625 + 225 = 4,450; the change is 2 × 10 × 5 = 100. (c) Current shares ignore potential competition; the harm lies in removing a future rival (killer acquisitions and the kill zone of Section 10.1), which a static index cannot see.

10.6. Case Studies and Debates

The materials below submit the chapter's regulatory architecture to empirical audit: the data-protection regime after its first years of enforcement, the gatekeeper regime in its opening phase, a small state that built its own digital infrastructure, and the frontier question the decade will answer.

Case 10.1. GDPR after the honeymoon: an empirical audit

Section 10.2 presented the General Data Protection Regulation as the foundation of the European model and the engine of the Brussels effect; the case's task is to weigh the accumulating evidence. On the regime's own terms the ledger shows genuine entries: a global diffusion of GDPR-style statutes across dozens of jurisdictions (Bradford, 2020); fines against major platforms scaling into the billions of euros; the entrenchment of rights - access, erasure, portability - that data subjects demonstrably exercise; and a compliance profession and privacy-by-design engineering culture that did not exist before the regulation created its demand. The critical ledger is equally documented. Enforcement runs through the supervisory authority of the member state where a firm establishes its European base, concentrating the docket of the largest platforms in a small jurisdiction whose resourcing became the system's acknowledged bottleneck - a design vulnerability of the one-stop-shop that reform proposals target. Consent, the regime's most visible instrument, degenerated at the interface into the banner fatigue every reader knows: a ritual of clicks that satisfies the form of autonomy while dark patterns steer its substance, illustrating Lessig's (1999) modalities - law commanded, code complied strategically.

The most instructive finding, however, concerns competition. Johnson et al. (2023), studying the web-technology market around the regulation's entry into force, documented an unintended consequence with this book's mechanisms written through it: compliance costs and legal risk fall relatively harder on small vendors and websites, and reliance shifted toward the largest platforms - whose consent infrastructure, legal departments, and first-party data absorb the regulation most easily - measurably increasing concentration in the affected markets. Privacy regulation, designed against the data extraction of Chapter 6, thereby strengthened, at the margin, the very gatekeepers targeted by the competition regime of Section 10.1 - the clearest available demonstration that the European architecture's pillars interact and can collide. Seminar questions: (1) Trace the mechanism from fixed compliance costs to concentration using Chapter 2's cost structures; which design changes - thresholds, standardized consent, public infrastructure - would decouple privacy from scale? (2) Is banner-consent failure a

failure of the law, the interface, or the concept of consent itself? Propose the alternative lawful basis you would rely on instead, and its costs. (3) Audit the Brussels effect against this case: what exactly diffused - the rights, the enforcement, or the paperwork?

Case 10.2. The gatekeeper regime's opening moves

The Digital Markets Act shifted platform regulation from litigation to obligation; its opening phase supplies a live experiment in whether ex ante rules can move markets that two decades of antitrust cases moved slowly. The early record exhibits the pattern regulatory theory predicts. Designations proceeded on the quantitative thresholds, sweeping in the expected constellation of operating systems, app stores, search, social, advertising, and intermediation services - with contestation at the margins over which services qualify. Compliance arrived in two registers: visible architectural change - choice screens, alternative app distribution and payment channels on mobile platforms, interoperability commitments in messaging, restrictions on combining data across services. The second register was what critics named compliance theatre: implementations formally satisfying an obligation while burdening the very behaviour it protects, as when alternative distribution was opened subject to new fees and friction whose net effect on developer economics was disputed, or when choice screens were designed with the defaults and orderings that behavioural research predicts will preserve incumbency. The Commission's first non-compliance investigations and penalty decisions, arriving within the regime's opening years, tested whether the obligations would be construed by their purpose or their letter.

2026 update. By mid-2026 the DMA had moved beyond designation and experimental compliance into a genuine enforcement phase. On 23 July 2026 the European Commission adopted two non-compliance decisions against Google, imposing fines of EUR 460 million for self-preferencing in Google Search and EUR 430 million for restrictions on steering users toward alternative purchase channels in Google Play. The analytical significance is larger than the monetary amount: the regime is now testing whether behavioural obligations can change ranking, distribution and payment architectures rather than merely alter formal terms. The case should therefore be taught as an evolving sequence - designation → implementation → specification → non-compliance finding → remedy - rather than as a one-off antitrust judgment (European Commission, 2026a).

The case belongs in a textbook precisely because it is unresolved: it allows the student to apply the book's mechanisms predictively rather than retrospectively. Seminar questions: (1) Select one DMA obligation and its contested implementation; using Chapters 3 and 5, argue

whether the implementation preserves the targeted market power, and design the test the regulator should run. (2) ‘Ex ante regulation trades error costs for speed.’ Audit the trade after the opening phase: which errors - over-inclusion, under-inclusion, ossification - have materialized? (3) The DMA regulates identified gatekeepers; AI foundation models are producing new candidate bottlenecks (Chapter 6). Draft the designation criteria you would apply to the model layer, and state what the DMA’s experience teaches about thresholds versus case-by-case designation.

Case 10.3. Estonia: the state as platform

Most of this chapter treats the state as a regulator of platforms built by others. Estonia shows the state as a platform builder. After regaining independence in 1991, a small country with limited administrative capacity chose to digitize public services rather than to replicate paper bureaucracy. Its core infrastructure is X-Road, a decentralized data-exchange layer launched in 2001 that lets public and private databases query one another securely, combined with a mandatory digital identity card introduced in 2002 that enables electronic signatures and authentication. On this foundation Estonia built internet voting in national elections from 2005 and, from 2014, e-Residency, which allows foreigners to obtain a digital identity in order to found and run an Estonian company online (Kattel & Mergel, 2019).

Read with this book’s mechanisms, X-Road is a two-sided platform whose sides are data holders and data users. The state governs it as the platform regulator of Section 3.7, but the design rests on interoperability rather than centralization (Section 3.8): data stay with their owners and are exchanged under logged, auditable rules - the institutional trust of Section 8.7 built into infrastructure. The ‘once-only’ principle, under which citizens should not be asked twice for information the state already holds, cuts the transaction costs of Chapter 2 for every public interaction. Kattel and Mergel (2019) caution, however, against treating the case as a simple success story. Its achievements depended on a particular historical window, a small and cohesive elite and a willingness to accept risks that larger states avoid. The 2007 cyber-attacks on Estonian institutions and a 2017 security flaw in its identity cards also showed how digital dependence creates new vulnerabilities - the resilience problem of Chapter 6. Seminar questions: (1) Is X-Road a public good, a club good or a platform? Use the taxonomy of Section 2.1 and the platform definition of Section 3.2. (2) Which elements of the Estonian model could a larger state adopt, and which depend on scale or history? (3) Compare Estonia’s public digital infrastructure with India’s UPI (Section 4.5a) and the proposed digital euro (Section 5.9) as strategies of digital sovereignty (Section 10.3).

Debate 10.1. Governing the frontier: should advanced AI development be licensed?

Motion: the development and deployment of frontier artificial intelligence models should require a licence, as nuclear operation, aviation, and pharmaceuticals do. *For the motion:* the scale economics of Chapter 6 have already concentrated frontier development among a handful of firms, so licensing formalizes an oligopoly that exists rather than creating one. The technology's error modes - from the classification harms of Section 6.2 to systemic risks in information (Chapter 9) and, prospectively, in security - are borne by third parties, the textbook condition for ex ante safety regulation rather than after-the-fact liability. Capability evaluations, incident reporting, and audit access are precisely the transparency that Pasquale's (2015) black box otherwise forecloses; and the alternative - racing dynamics among firms and states - is the reward structure Baumol (1990) warns about, steering talent toward speed over safety. *Against the motion:* licensing frontier development entrenches exactly the incumbents whose power Chapters 3, 6, and 10 problematized - the compliance moat of Case 10.1 rebuilt one layer up - while open-source and academic development, the historical engine of the field's progress and its scrutiny (Chapter 2's commons), is burdened or criminalized. The harms invoked are heterogeneous, and the demonstrated ones (bias, misinformation) are governed better by domain rules - the DSA's systemic-risk audits, discrimination law, product liability - than by gating the general technology. Regulatory capacity to evaluate frontier capabilities lags the frontier by construction, making the licence either a rubber stamp or a freeze. Finally, the geopolitics of Bradford (2023) ensure that unilateral licensing relocates rather than restrains development. The debate should be run twice - once assuming capabilities plateau, once assuming they compound - and each side must state which assumption its position secretly requires, and what evidence within five years would make it switch sides.

10.7. In Depth: Governing Intelligence - Comparative Architectures of AI Regulation

Section 10.2 mapped the European model for data and platforms; this advanced section extends the comparison to the technology of Chapter 6, because AI regulation is where Bradford's (2023) three empires are composing their next constitutions in real time, and where every analytical thread of the book - classification harms, scale economics, systemic risk, the Brussels effect's reach - converges on live institutional design.

The rights-driven architecture: risk tiers and their logic

The European Union legislated first and most comprehensively. The Artificial Intelligence Act (Regulation (EU) 2024/1689) organizes the entire technology through a *risk pyramid*: prohibited practices at the apex - social scoring by public or private actors (Case 8.2 pre-empted by statute), certain manipulative and biometric-categorization uses; a *high-risk* tier - AI in employment, credit, education, essential services, law enforcement, the classification situations of Chapters 6 and 8 - subject to ex ante conformity assessment, data-governance and accuracy requirements, human oversight, and logging; a transparency tier for systems that interact with people or generate content; and minimal obligations for the remainder. Onto this use-based pyramid the legislative process grafted, mid-drafting, the model layer the pyramid had not anticipated: obligations for *general-purpose AI models*, scaled by a compute threshold above which systemic-risk duties - evaluations, adversarial testing, incident reporting - attach to the frontier systems of Section 6.7. The architecture's logic is continuous with the GDPR lineage: rights protected ex ante, obligations graduated by risk, enforcement by fine scaled to turnover - and so are its predictable frictions, rehearsed in Case 10.1: conformity assessment as compliance moat favouring scale, definitional boundary wars over what counts as high-risk, and the standing question whether the Brussels effect extends to a technology whose supply chain, unlike a website's cookie banner, can be regionally partitioned.

2026 implementation update. The AI Act is no longer only a legislative architecture. From 2 August 2026 the European Commission's AI Office and national competent authorities began enforcing the provisions that had become applicable, including obligations for providers of general-purpose AI models and transparency requirements for specified AI systems and synthetic content. The timetable remains staged, and was stretched further by the AI Omnibus (Regulation (EU) 2026/1744), in force from 27 July 2026: obligations for stand-alone high-risk systems (Annex III) apply from 2 December 2027 and those for AI embedded in

regulated products (Annex I) from 2 August 2028, so students should distinguish adoption of the regulation, legal applicability of particular obligations, and operational enforcement capacity. This staged implementation is itself an economic variable because firms face compliance costs at different points in the value stack and on different time horizons. (European Commission, 2026b; European Parliament & Council of the European Union, 2026)

The market-driven and state-driven alternatives

The United States, true to Bradford's (2023) taxonomy, produced not a statute but a lattice: executive-branch direction to agencies (a 2023 executive order on AI, revoked in January 2025 in favour of a deregulatory, competitiveness-first agenda), voluntary safety-testing and reporting commitments negotiated with frontier developers, sectoral enforcement through existing law - the discrimination, consumer-protection, and product-liability regimes that Debate 10.1's opponents preferred - and vigorous sub-federal experimentation, with states legislating on deepfakes, automated decision notices, and training-data disclosure. The architecture's strengths and weaknesses are the mirror of Europe's: adaptive and innovation-friendly by design, it is also revocable by election (as the 2025 reversal demonstrated), fragmented across jurisdictions, and dependent on the voluntary cooperation of the regulated - governance by memorandum where the EU governs by regulation. China's state-driven architecture, meanwhile, regulated earliest and most instrumentally: binding rules on recommendation algorithms, deep synthesis, and generative services arrived years before either Western framework, combining genuine consumer-protection and labelling provisions with requirements - content alignment with core values, security assessments, algorithm registries filed with the regulator - that subordinate the technology to information control, the sovereignty of Section 10.3 in its authoritarian register. The registry instrument deserves the student's particular attention: a state that requires algorithms to be disclosed to it, while their subjects receive only labels, has resolved Pasquale's (2015) black-box problem asymmetrically - transparency upward, opacity downward.

The convergence problem and the analytical payoff

Between the architectures runs the coordination question the decade will answer. The technology's supply chain is global (Section 6.7: chips, cloud, models, applications cross all three empires), its risks travel (Chapter 9's information disorder respects no jurisdiction), and firms face the compliance-divisibility calculus of the Brussels effect - build one governance

stack to the strictest standard, or partition products by region, a choice model providers are visibly making differently for different obligations. International coordination has produced summit declarations, safety-institute networks, and standards-body activity - the infrastructure governance of DeNardis (2014) migrating to model evaluations - but no instrument with the bite of trade law, and the geopolitical overlay (export controls on the chip layer, explicitly framed as strategic competition) ensures that AI governance is also industrial policy conducted by other means. For the student, the section's payoff is a template rather than a forecast: any AI-governance proposal can now be located on three axes - *what* is regulated (uses, models, or compute), *when* (ex ante conformity or ex post liability), and *for whom* (rights of the classified, safety of the public, or security of the state). The book's mechanisms predict each choice's characteristic failure: use-based rules chase a moving frontier, model-based rules build compliance moats (Case 10.1), compute thresholds invite efficient evasion (Section 6.7's algorithmic progress), liability waits for the harm, and conformity assessment freezes the very iteration that Chapter 4's economics of experimentation showed to be the sector's engine. Discussion questions: (1) Place the three architectures on the three axes, and identify for each the failure mode the book's mechanisms predict. (2) The AI Act prohibits social scoring by public and private actors alike, yet leaves much of the private scoring of Chapter 8 - platform ratings, customer values - outside its prohibitions and treats credit scoring as high-risk rather than banned; defend and attack where these lines fall. (3) Will the Brussels effect operate for the AI Act? Work through the divisibility calculus for one concrete obligation. (4) Draft the one coordination instrument you would prioritize internationally - evaluation standards, incident reporting, or compute accounting - and explain what each empire would demand in exchange.

Instrument	Primary object	Core mechanism	Characteristic trade-off
GDPR	Personal data processing	Rights, lawful bases, accountability	Privacy protection vs fixed compliance costs and consent fatigue
DSA	Intermediaries and very large platforms	Due diligence, transparency, systemic-risk duties	Public-interest obligations vs over-removal and audit complexity
DMA	Designated gatekeepers	Ex ante contestability obligations	Speed and interoperability vs ossification and compliance engineering
AI Act	AI systems and GPAI models	Risk tiers, transparency, model obligations	Safety/rights vs entry barriers and fast-moving definitions

Table 10.1. European digital regulation as a portfolio of partially overlapping instruments.

Conclusion: The Digital Economy as a Constructed Order

This book opened with a claim and closes having, its authors hope, earned it: the digital economy is simultaneously an economic mechanism, a social institution, and a media order, and it yields its intelligibility only to an analysis willing to hold all three in view. The economic mechanism came first because it is foundational: the peculiar properties of information - non-rival, cheap to copy, costly to create - generate the pricing strategies, the network dynamics, and the platform structures from which everything else follows, and the reader who has internalized Chapters 2 and 3 possesses a compact engine that explains, without mystery, the zero price and the walled garden, the colossal bundle and the subsidized ride, the tipped market and the enveloped rival. Economics, this book has insisted, is not optional equipment for the critic of digital capitalism; it is the grammar in which the object is written.

But the mechanism, left alone, explains a world that does not exist - a frictionless space of atomized agents - and the book's second movement supplied what the mechanism presupposes. Markets among strangers required engineered trust, and the engineering imported the social structure of its users; opportunity travelled through networks that are socially closed; participation presupposed access, skills, and dispositions that are stratified; and the classifications by which the data economy sees - creditworthy, relevant, risky, visible - returned to their subjects as allocated fates. The sociology of Chapters 4, 7, and 8 is therefore not a humanistic supplement to the economics but its completion: embeddedness, stratification, and

quantification are where the mechanism touches ground, and where its costs and benefits find their unequal addresses.

The third movement followed the object into the domain where its consequences are most public. The dominant firms of the digital economy are, by revenue logic, attention merchants; their feeds perform, algorithmically and for profit, the gatekeeping that constitutes publics; and the collapse of journalism's accidental subsidy, the selection bias of engagement, and the constitutive character of moderation together transformed the infrastructure of democratic communication while its new custodians disclaimed the office. The media lens, like the sociological one, proved to be not a perspective on the digital economy but a description of part of what it is.

Three theses may serve as the book's summary. First, *the digital economy is layered and cumulative*: five waves, each presupposing its predecessors, have deposited an order in which chips, networks, platforms, data, and models stack - and in which power accrues disproportionately at the layers of intermediation and infrastructure. Second, *its markets are designed*: reputation systems, price structures, rankings, and default settings are institutional choices with distributional consequences, made mostly by private actors and revisable - which is why the book has treated design details as the constitutional law of digital markets. Third, *its future is institutional, not technological*: the same machinery of prediction and connection supports surveillance or privacy, so-so automation or augmentation, enclosure or commons, fragmentation or public reason - and the selection among these is performed by the reward structures, regulations, and collective choices that Baumol, Lessig, and Acemoglu and Johnson, in their different idioms, all identified as the true variables.

For the three readers to whom the book is addressed, three departures suggest themselves. The economist leaves with an enlarged object: market design that allocates life-chances, and competition that occurs between constitutional orders as much as between firms. The sociologist leaves with an enlarged toolkit: the cost structures and network dynamics without which platform power is a black box, and with which it is a mechanism. The media scholar leaves with an enlarged jurisdiction: the political economy of communication now runs through auction theory and antitrust, and the public sphere has a balance sheet. And all three leave, the authors hope, with the book's methodological habit: to ask of every digital arrangement - every price of zero, every score, every feed - the questions that have organized these chapters. Who pays, and who is subsidized? What is measured, and what does the measure discipline? Where does the risk sit, and where the power? Ask them systematically, and the digital economy ceases to be a spectacle of disruption and becomes what this book has tried to

make it: an intelligible, constructed, and therefore contestable order - the proper object of scholarship, and of citizenship.

Epilogue: Reading the Next Wave

Textbooks on fast-moving subjects age in a characteristic way: their examples date first, their empirical claims second, and their frameworks - if the frameworks were any good - last of all. This epilogue is written for the reader who opens the book after its examples have dated, and it consists of a single exercise: a checklist for reading whatever wave is cresting at the time of reading, assembled entirely from the machinery of the preceding chapters. If the book has succeeded, the checklist will feel obvious; that feeling is the book's intended residue.

First, locate the wave's *cost collapse*. Every wave in this book's chronology was, at bottom, the collapse of a cost: computation (the first and second waves), distribution (the third), coordination (the fourth), prediction and now generation (the fifth). Ask of the new technology: which of the costs in Goldfarb and Tucker's catalogue - or which cost they did not need to name - has fallen by an order of magnitude? The answer predicts, before any product exists, where activity will reorganize: the newly cheap thing will be used lavishly, its complements will appreciate, and its former economizers - the occupations, firms, and institutions built to husband the once-scarce resource - will face the choice between redeployment and obsolescence that Chapters 6 and 7 anatomized.

Second, find the *bottleneck that remains*. Abundance is never general; it relocates scarcity. When information became abundant, attention became the commodity (Chapter 9); when computation became rentable, data and distribution became the moats (Chapters 5 and 6); when models became capable, compute, energy, and trusted deployment became the constraints (Section 6.7). The durable market power of the coming wave will sit at its remaining bottleneck, not at its headline capability - and the firms to watch are those quietly assembling positions there while public attention prices the abundance.

Third, audit the *institutional settlement in formation*. Each wave, this book has argued, is stabilized by settlements - of property (Chapter 2 and Case 6.2), of classification (Chapter 7's employment struggle), of liability and access (Chapters 9 and 10) - that are negotiated early, in the fog, and then govern for decades: the software licence, the intermediary immunity, the consent regime. Identify the settlement currently under negotiation for the new wave, ask who is at the table and who is arriving late, and remember the book's most practical lesson: the parties who treat the settlement as a technical detail will discover, a decade on, that it was the constitution.

Fourth, run the *three lenses in order, then in reverse*. Begin as the economist: what are the cost structures, the network effects, the price structures - who pays, who is subsidized, what

is measured, where does the risk sit? Continue as the sociologist: through which networks does access run, which classifications are being built, whose participation is presupposed and whose is foreclosed? Finish as the media scholar: what does the new machinery do to the flow of public information, and who curates it under what accountability? Then reverse the order and notice what the first pass missed - for the book's deepest methodological claim is that each lens sees what the others cannot, and the object, which was never divided, requires all three.

The checklist ends, as the book did, on the point that distinguishes analysis from spectatorship. The waves are not weather. Their cost collapses are physics and engineering, but their bottlenecks are partly built, their settlements are wholly chosen, and their distributional outcomes - who gains, who bears the transition, what survives of the commons and the public sphere - are the running total of decisions by firms, states, workers, and publics, most of them taken while the wave still looked like someone else's subject. The reader who has finished this book is, by that fact, one of the people at the table. The authors' closing hope is immodest in the only way a textbook is permitted to be: that the reader will find the table, and use the checklist there.

Capstone Project: Audit a Digital Market

Select one contemporary digital market, venture or technology and analyse it with the full apparatus of the book. Your report should identify the cost collapse that made the market possible; the remaining bottleneck; the relevant network effects and switching costs; the business model and unit economics; the distribution of risk among firm, worker, user and complementor; the role of data and algorithmic classification; effects on attention or public communication where relevant; and the institutional settlement governing competition, rights and accountability. Conclude not with a prediction but with two contrasting scenarios and the observable indicators that would distinguish them.

Assessment criterion: a strong report makes mechanisms explicit, distinguishes evidence from assumption, and shows where the three lenses - economics, sociology and media studies - produce different but complementary diagnoses of the same object.

Glossary of Core Concepts

Algorithmic curation: The automated selection, ranking and recommendation of content that determines what users see in feeds, search results and recommendation surfaces.

Algorithmic management: The delegation of direction, evaluation and discipline of workers to software-mediated systems.

Attention economy: An economic order in which scarce human attention is allocated, measured and monetized.

Brussels effect: The process by which EU rules become de facto global standards because firms find uniform compliance cheaper than differentiating by market.

CAC (customer acquisition cost): Average cost of acquiring a new paying customer through sales and marketing activity.

Central bank digital currency (CBDC): A digital form of money issued directly by a central bank and usable by the public, as distinct from bank deposits and private digital tokens.

Churn: The proportion of customers or recurring revenue lost during a specified period.

Complementor: A firm or producer whose offering increases the value of a platform while depending on the platform for access.

Data colonialism: A critical account of datafication as the appropriation of human life through data, likened to the historical colonial appropriation of land and resources.

Datafication: The conversion of activities, relations or events into machine-readable data.

Digital divide: Inequality in motivation, access, skills, patterns of use and outcomes from digital technologies.

Digital entrepreneurship: Opportunity pursuit in which digital technologies are constitutive of the venture product or process.

Embeddedness: The principle that economic action is shaped by, and depends on, ongoing networks of social relations, norms and institutions.

Envelopment: Entry by a platform into an adjacent market by bundling functionality and leveraging an existing user base.

Experience good: A good whose quality or value is difficult to know before consumption.

Freemium: A versioning model in which a functional free tier recruits users and a minority converts to paid tiers.

Gatekeeper: Under the Digital Markets Act, a designated provider of core platform services that serves as an important gateway between business users and end users and enjoys an entrenched and durable position.

General purpose technology: A pervasive, improving technology that generates complementary innovation across many sectors.

Indirect network effect: An increase in value to one user group caused by growth of another group on the same platform.

Interoperability: The capacity of independently governed systems to exchange data or enable cross-system interaction.

Lock-in: A condition in which switching costs, network effects or accumulated data make exit economically difficult.

LTV (lifetime value): Expected contribution margin generated by a customer over the duration of the relationship.

Multi-homing: Participation by a user or complementor in more than one competing platform.

Non-rivalry: A property of goods for which one person's use does not materially reduce availability to others.

Open banking: Regulation requiring banks to let licensed third parties, with the customer's consent, access account data and initiate payments; a form of mandated data portability.

Platformization: The penetration of platform infrastructures, governance and metrics into sectors of economic and cultural production.

Presumption of employment: A legal rule under which a person working through a platform is presumed to be an employee unless the platform proves otherwise, as in the EU Platform Work Directive.

Product-market fit: A condition in which a product repeatedly satisfies a sufficiently strong market demand, visible in sustained adoption and retention.

Runway: The period a venture can continue operating at its present net burn rate before cash is exhausted.

Surveillance capitalism: Zuboff's term for an accumulation logic in which behavioural data are extracted beyond service needs and converted into prediction products sold to third parties.

Switching cost: Any monetary, technical, informational or social cost incurred when moving from one provider or platform to another.

Tipping: A self-reinforcing shift by which a network market converges toward one or a few dominant providers.

Two-sided market: A market in which an intermediary coordinates distinct user groups whose participation creates cross-side effects.

Versioning: Designing differentiated variants of a low-marginal-cost product to induce self-selection among customers with different willingness to pay.

Winner-take-most: A market structure in which feedback effects produce a highly skewed distribution of users, revenue or attention without necessarily creating a literal monopoly.

Subject Guide: Core Topics Across the Book

This guide is designed as a teaching and revision index. It points readers to the chapters in which a concept is introduced, formalized, or revisited through another disciplinary lens.

Topic	Chapters	Primary analytical use
AI / foundation models	6, 7, 9, 10	production economics; labour; media; regulation
Algorithmic management	7, 8	labour control; quantification and trust
Attention economy	5, 9	business model; media political economy
Business models	4, 5	venture design; value capture
Cybersecurity and resilience	6, 10	externalities; infrastructure; governance
Data and privacy	2, 6, 10	information economics; extraction; rights
Digital divide	1, 8	measurement; stratified participation
Digital goods	1, 2, 5	cost structure; pricing; monetization
Digital labour	4, 7, 8	entrepreneurship; task markets; inequality
Digital sovereignty	6, 10	infrastructure; governance
Entrepreneurial ecosystems	4	networks; finance; institutions
Generative search / AI answers	6, 9	AI production; media gateways

Interoperability	3, 10	competition mechanism; regulatory remedy
Network effects	1, 3, 5	increasing returns; platforms; growth
Platform governance	3, 8, 9, 10	private ordering; reputation; media; law
Product-market fit	4, 5	experimentation; unit economics
Reputation and trust	2, 3, 8	verification; platform exchange; sociology
SaaS / subscriptions	5	pricing; retention; LTV/CAC
Surveillance / tracking	2, 6, 9	cost collapse; data extraction; attention
Venture finance	4, 5	staging; growth economics

References

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy, 40*(121), 13-58. https://doi.org/10.1093/epolic/eiae042

Acemoglu, D., & Johnson, S. (2023). *Power and progress: Our thousand-year struggle over technology and prosperity*. PublicAffairs.

Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives, 33*(2), 3-30. https://doi.org/10.1257/jep.33.2.3

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature, 54*(2), 442-492. https://doi.org/10.1257/jel.54.2.442

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics, 84*(3), 488-500. https://doi.org/10.2307/1879431

Allcott, H., Braghieri, L., Eichmeyer, S., & Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review, 110*(3), 629-676. https://doi.org/10.1257/aer.20190658

Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives, 31*(2), 211-236. https://doi.org/10.1257/jep.31.2.211

Anderson, C. (2006). *The long tail: Why the future of business is selling less of more*. Hyperion.

Anderson, R., & Moore, T. (2006). The economics of information security. *Science, 314*(5799), 610-613. https://doi.org/10.1126/science.1130992

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). *Machine bias*. ProPublica. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM, 53*(4), 50-58. https://doi.org/10.1145/1721654.1721672

Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics, 37*(3), 668-691. https://doi.org/10.1111/j.1756-2171.2006.tb00037.x

Arntz, M., Gregory, T., & Zierahn, U. (2016). *The risk of automation for jobs in OECD countries: A comparative analysis* (OECD Social, Employment and Migration Working Papers No. 189). OECD Publishing. https://doi.org/10.1787/5jlz9h56dvq7-en

Arrieta-Ibarra, I., Goff, L., Jiménez-Hernández, D., Lanier, J., & Weyl, E. G. (2018). Should we treat data as labor? Moving beyond "free". *AEA Papers and Proceedings, 108*, 38-42. https://doi.org/10.1257/pandp.20181003

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In *The rate and direction of inventive activity: Economic and social factors* (pp. 609-626). Princeton University Press.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal, 99*(394), 116-131. https://doi.org/10.2307/2234208

Auer, R., Frost, J., Gambacorta, L., Monnet, C., Rice, T., & Shin, H. S. (2022). Central bank digital currencies: Motives, economic implications, and the research frontier. *Annual Review of Economics, 14*, 697-721. https://doi.org/10.1146/annurev-economics-051420-020324

Autio, E., Nambisan, S., Thomas, L. D. W., & Wright, M. (2018). Digital affordances, spatial affordances, and the genesis of entrepreneurial ecosystems. *Strategic Entrepreneurship Journal, 12*(1), 72-95. https://doi.org/10.1002/sej.1266

Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives, 29*(3), 3-30. https://doi.org/10.1257/jep.29.3.3

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics, 118*(4), 1279-1333. https://doi.org/10.1162/003355303322552801

Babina, T., Bahaj, S., Buchak, G., De Marco, F., Foulis, A., Gornall, W., Mazzola, F., & Yu, T. (2025). Customer data access and fintech entry: Early evidence from open banking. *Journal of Financial Economics, 169*, Article 103950. https://doi.org/10.1016/j.jfineco.2024.103950

Bakos, Y., & Brynjolfsson, E. (1999). Bundling information goods: Pricing, profits, and efficiency. *Management Science, 45*(12), 1613-1630. https://doi.org/10.1287/mnsc.45.12.1613

Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science, 348*(6239), 1130-1132. https://doi.org/10.1126/science.aaa1160

Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review, 104*(3), 671-732. https://doi.org/10.15779/Z38BG31

Barrero, J. M., Bloom, N., & Davis, S. J. (2023). The evolution of work from home. *Journal of Economic Perspectives, 37*(4), 23-49. https://doi.org/10.1257/jep.37.4.23

Bartz v. Anthropic PBC, No. 3:24-cv-05417-WHA (N.D. Cal. June 23, 2025).

Baumol, W. J. (1990). Entrepreneurship: Productive, unproductive, and destructive. *Journal of Political Economy, 98*(5), 893-921. https://doi.org/10.1086/261712

Bell, D. (1973). *The coming of post-industrial society: A venture in social forecasting*. Basic Books.

Benkler, Y. (2006). *The wealth of networks: How social production transforms markets and freedom*. Yale University Press.

Berg, J., Furrer, M., Harmon, E., Rani, U., & Silberman, M. S. (2018). *Digital labour platforms and the future of work: Towards decent work in the online world*. International Labour Office.

Blake, T., Nosko, C., & Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica, 83*(1), 155-190. https://doi.org/10.3982/ECTA12423

Blank, S. (2013). Why the lean start-up changes everything. *Harvard Business Review, 91*(5), 63-72.

Bloom, N., Liang, J., Roberts, J., & Ying, Z. J. (2015). Does working from home work? Evidence from a Chinese experiment. *The Quarterly Journal of Economics, 130*(1), 165-218. https://doi.org/10.1093/qje/qju032

Boldrin, M., & Levine, D. K. (2008). *Against intellectual monopoly*. Cambridge University Press. https://doi.org/10.1017/CBO9780511510854

Bork, R. H. (1978). *The antitrust paradox: A policy at war with itself*. Basic Books.

Boudreau, K. J. (2010). Open platform strategies and innovation: Granting access vs. devolving control. *Management Science, 56*(10), 1849-1872. https://doi.org/10.1287/mnsc.1100.1215

Boudreau, K. J., & Hagiu, A. (2009). Platform rules: Multi-sided platforms as regulators. In A. Gawer (Ed.), *Platforms, markets and innovation* (pp. 163-191). Edward Elgar. https://doi.org/10.4337/9781849803311.00013

Bradford, A. (2020). *The Brussels effect: How the European Union rules the world*. Oxford University Press. https://doi.org/10.1093/oso/9780190088583.001.0001

Bradford, A. (2023). *Digital empires: The global battle to regulate technology*. Oxford University Press. https://doi.org/10.1093/oso/9780197649268.001.0001

Braghieri, L., Levy, R., & Makarin, A. (2022). Social media and mental health. *American Economic Review, 112*(11), 3660-3693. https://doi.org/10.1257/aer.20211218

Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies: “Engines of growth”? *Journal of Econometrics, 65*(1), 83-108. https://doi.org/10.1016/0304-4076(94)01598-T

Brynjolfsson, E. (1993). The productivity paradox of information technology. *Communications of the ACM, 36*(12), 66-77. https://doi.org/10.1145/163298.163309

Brynjolfsson, E., Collis, A., Diewert, W. E., Eggers, F., & Fox, K. J. (2019). *GDP-B: Accounting for the value of new and free goods in the digital economy* (NBER Working Paper No. 25695). National Bureau of Economic Research. https://doi.org/10.3386/w25695

Brynjolfsson, E., Collis, A., & Eggers, F. (2019). Using massive online choice experiments to measure changes in well-being. *Proceedings of the National Academy of Sciences, 116*(15), 7250-7255. https://doi.org/10.1073/pnas.1815663116

Brynjolfsson, E., Hu, Y., & Smith, M. D. (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science, 49*(11), 1580-1596. https://doi.org/10.1287/mnsc.49.11.1580.20580

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889-942. https://doi.org/10.1093/qje/qjae044

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton.

Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 23-57). University of Chicago Press. https://doi.org/10.7208/chicago/9780226613475.003.0001

Brynjolfsson, E., & Smith, M. D. (2000). Frictionless commerce? A comparison of Internet and conventional retailers. *Management Science, 46*(4), 563-585. https://doi.org/10.1287/mnsc.46.4.563.12061

Bukht, R., & Heeks, R. (2017). *Defining, conceptualising and measuring the digital economy* (Development Informatics Working Paper No. 68). University of Manchester, Global Development Institute. https://doi.org/10.2139/ssrn.3431732

Cabral, L., Haucap, J., Parker, G., Petropoulos, G., Valletti, T., & Van Alstyne, M. (2021). *The EU Digital Markets Act: A report from a panel of economic experts*. Publications Office of the European Union. https://doi.org/10.2760/139337

Caillaud, B., & Jullien, B. (2003). Chicken & egg: Competition among intermediation service providers. *The RAND Journal of Economics, 34*(2), 309-328. https://doi.org/10.2307/1593720

Castells, M. (2010). *The rise of the network society* (2nd ed.). Wiley-Blackwell. (Original work published 1996)

Castronova, E. (2005). *Synthetic worlds: The business and culture of online games*. University of Chicago Press.

Chan, A., Salganik, R., Markelius, A., Pang, C., Rajkumar, N., Krasheninnikov, D., Langosco, L., He, Z., Duan, Y., Carroll, M., et al. (2023). Harms from increasingly agentic algorithmic systems. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM. https://doi.org/10.1145/3593013.3594033

Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data, 5*(2), 153-163. https://doi.org/10.1089/big.2016.0047

Christin, A. (2020). *Metrics at work: Journalism and the contested meaning of algorithms*. Princeton University Press. https://doi.org/10.1515/9780691200002

Cook, C., Diamond, R., Hall, J. V., List, J. A., & Oyer, P. (2021). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies, 88*(5), 2210-2238. https://doi.org/10.1093/restud/rdaa081

Couldry, N., & Mejias, U. A. (2019). Data colonialism: Rethinking big data's relation to the contemporary subject. *Television & New Media, 20*(4), 336-349. https://doi.org/10.1177/1527476418796632

Coyle, D. (2014). *GDP: A brief but affectionate history*. Princeton University Press.

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Crémer, J., de Montjoye, Y.-A., & Schweitzer, H. (2019). *Competition policy for the digital era*. Publications Office of the European Union. https://doi.org/10.2763/407537

Cui, R., Li, J., & Zhang, D. J. (2020). Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb. *Management Science, 66*(3), 1071-1094. https://doi.org/10.1287/mnsc.2018.3273

Cunningham, C., Ederer, F., & Ma, S. (2021). Killer acquisitions. *Journal of Political Economy, 129*(3), 649-702. https://doi.org/10.1086/712506

Cunningham, S., & Craig, D. (2019). *Social media entertainment: The emergence of a new screen ecology*. New York University Press.

Cusumano, M. A., Gawer, A., & Yoffie, D. B. (2019). *The business of platforms: Strategy in the age of digital competition, innovation, and power*. Harper Business.

David, P. A. (1985). Clio and the economics of QWERTY. *The American Economic Review, 75*(2), 332-337.

David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *The American Economic Review, 80*(2), 355-361.

De Stefano, V. (2016). The rise of the "just-in-time workforce": On-demand work, crowdwork, and labor protection in the gig economy. *Comparative Labor Law & Policy Journal, 37*(3), 471-504.

DeNardis, L. (2014). *The global war for Internet governance*. Yale University Press. https://doi.org/10.12987/yale/9780300181357.001.0001

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory, 3*(2), 156-168. https://doi.org/10.1016/0022-0531(71)90013-5

DiMaggio, P., Hargittai, E., Neuman, W. R., & Robinson, J. P. (2001). Social implications of the Internet. *Annual Review of Sociology, 27*, 307-336. https://doi.org/10.1146/annurev.soc.27.1.307

Dranove, D., & Jin, G. Z. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature, 48*(4), 935-963. https://doi.org/10.1257/jel.48.4.935

Dubal, V. B. (2017). Wage slave or entrepreneur? Contesting the dualism of legal worker identities. *California Law Review, 105*(1), 65-123. https://doi.org/10.15779/Z38M84Z

Duffy, B. E. (2017). *(Not) getting paid to do what you love: Gender, social media, and aspirational work*. Yale University Press.

Edelman, B., Luca, M., & Svirsky, D. (2017). Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics, 9*(2), 1-22. https://doi.org/10.1257/app.20160213

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review, 97*(1), 242-259. https://doi.org/10.1257/aer.97.1.242

Eghbal, N. (2020). *Working in public: The making and maintenance of open source software*. Stripe Press.

Einav, L., Farronato, C., & Levin, J. (2016). Peer-to-peer markets. *Annual Review of Economics, 8*, 615-635. https://doi.org/10.1146/annurev-economics-080315-015334

W. Eisenmann, T., Parker, G., & Van Alstyne, M. W. (2006). Strategies for two-sided markets. *Harvard Business Review, 84*(10), 92-101.

Eisenmann, T., Parker, G., & Van Alstyne, M. W. (2011). Platform envelopment. *Strategic Management Journal, 32*(12), 1270-1285. https://doi.org/10.1002/smj.935

Ellison, G., & Ellison, S. F. (2009). Search, obfuscation, and price elasticity on the Internet. *Econometrica, 77*(2), 427-452. https://doi.org/10.3982/ECTA5708

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science, 384*(6702), 1306-1308. https://doi.org/10.1126/science.adj0998

Epic Games, Inc. v. Apple, Inc., 67 F.4th 946 (9th Cir. 2023).

Epic Games, Inc. v. Apple Inc., No. 4:20-cv-05640-YGR (N.D. Cal. Apr. 30, 2025).

Espeland, W. N., & Stevens, M. L. (1998). Commensuration as a social process. *Annual Review of Sociology, 24*, 313-343. https://doi.org/10.1146/annurev.soc.24.1.313

European Commission. (2023). *Proposal for a regulation of the European Parliament and of the Council on the establishment of the digital euro* (COM(2023) 369 final).

European Commission. (2026a, July 23). Commission fines Google EUR 890 million for breaches of the Digital Markets Act. Digital Markets Act. https://digital-markets-act.ec.europa.eu/commission-fines-google-eur890-million-breaches-digital-markets-act-2026-07-23_en

European Commission. (2026b, July 31). Commission starts enforcing AI Act rules and new transparency requirements on 2 August. Shaping Europe's Digital Future. https://digital-strategy.ec.europa.eu/en/news/commission-starts-enforcing-ai-act-rules-and-new-transparency-requirements-2-august

European Commission. (2026c). Directive on measures for a high common level of cybersecurity across the Union (NIS2 Directive). Shaping Europe's Digital Future. Updated July 2, 2026. https://digital-strategy.ec.europa.eu/en/policies/nis2-directive

European Commission v Hungary, Case C-596/19 P (Court of Justice of the European Union, 16 March 2021).

European Parliament & Council of the European Union. (2015). Directive (EU) 2015/2366 on payment services in the internal market (PSD2). *Official Journal of the European Union*, L 337, 35-127.

European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union*, L 119, 1-88.

European Parliament & Council of the European Union. (2022a). Regulation (EU) 2022/1925 (Digital Markets Act). *Official Journal of the European Union*, L 265, 1-66.

European Parliament & Council of the European Union. (2022b). Regulation (EU) 2022/2065 (Digital Services Act). *Official Journal of the European Union*, L 277, 1-102.

European Parliament & Council of the European Union. (2022c). Regulation (EU) 2022/868 (Data Governance Act). *Official Journal of the European Union*, L 152, 1-44.

European Parliament & Council of the European Union. (2023a). Regulation (EU) 2023/1114 on markets in crypto-assets (MiCA). *Official Journal of the European Union*, L 150, 40-205.

European Parliament & Council of the European Union. (2023b). Regulation (EU) 2023/2854 (Data Act). *Official Journal of the European Union*, L series.

European Parliament & Council of the European Union. (2024a). Directive (EU) 2024/2831 on improving working conditions in platform work. *Official Journal of the European Union*, L series.

European Parliament & Council of the European Union. (2024b). Regulation (EU) 2024/1689 (Artificial Intelligence Act). *Official Journal of the European Union*, L series.

European Parliament & Council of the European Union. (2026). Regulation (EU) 2026/1744 (Digital Omnibus on AI, amending Regulation (EU) 2024/1689). *Official Journal of the European Union*, L series.

Evans, D. S. (2009). The online advertising industry: Economics, evolution, and privacy. *Journal of Economic Perspectives, 23*(3), 37-60. https://doi.org/10.1257/jep.23.3.37

Evans, D. S., & Schmalensee, R. (2016). *Matchmakers: The new economics of multisided platforms*. Harvard Business Review Press.

Ewens, M., & Townsend, R. R. (2020). Are early stage investors biased against women? *Journal of Financial Economics, 135*(3), 653-677. https://doi.org/10.1016/j.jfineco.2019.07.002

Farrell, J., & Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. In M. Armstrong & R. Porter (Eds.), *Handbook of industrial organization* (Vol. 3, pp. 1967-2072). North-Holland. https://doi.org/10.1016/S1573-448X(06)03031-7

Farrell, J., & Saloner, G. (1985). Standardization, compatibility, and innovation. *The RAND Journal of Economics, 16*(1), 70-83. https://doi.org/10.2307/2555589

Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly, 80*(S1), 298-320. https://doi.org/10.1093/poq/nfw006

Fleder, D., & Hosanagar, K. (2009). Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management Science, 55*(5), 697-712. https://doi.org/10.1287/mnsc.1080.0974

Fourcade, M., & Healy, K. (2017). Seeing like a market. *Socio-Economic Review, 15*(1), 9-29. https://doi.org/10.1093/ser/mww033

Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S., & Friday, A. (2021). The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns, 2*(9), Article 100340. https://doi.org/10.1016/j.patter.2021.100340

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114*, 254-280. https://doi.org/10.1016/j.techfore.2016.08.019

Furman, J., Coyle, D., Fletcher, A., McAuley, D., & Marsden, P. (2019). *Unlocking digital competition: Report of the Digital Competition Expert Panel*. HM Treasury.

Gawer, A., & Cusumano, M. A. (2014). Industry platforms and ecosystem innovation. *Journal of Product Innovation Management, 31*(3), 417-433. https://doi.org/10.1111/jpim.12105

Gillespie, T. (2018). *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.

E. Goldfarb, A., & Tucker, C. E. (2011). Privacy regulation and online advertising. *Management Science, 57*(1), 57-71. https://doi.org/10.1287/mnsc.1100.1246

Goldfarb, A., & Tucker, C. E. (2019). Digital economics. *Journal of Economic Literature, 57*(1), 3-43. https://doi.org/10.1257/jel.20171452

Goldhaber, M. H. (1997). The attention economy and the Net. *First Monday, 2*(4). https://doi.org/10.5210/fm.v2i4.519

Gompers, P., & Lerner, J. (2001). The venture capital revolution. *Journal of Economic Perspectives, 15*(2), 145-168. https://doi.org/10.1257/jep.15.2.145

Goos, M., Manning, A., & Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review, 104*(8), 2509-2526. https://doi.org/10.1257/aer.104.8.2509

Graham, M., & Anwar, M. A. (2019). The global gig economy: Towards a planetary labour market? *First Monday, 24*(4). https://doi.org/10.5210/fm.v24i4.9913

Graham, M., Hjorth, I., & Lehdonvirta, V. (2017). Digital labour and development: Impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: European Review of Labour and Research, 23*(2), 135-162. https://doi.org/10.1177/1024258916687250

Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology, 91*(3), 481-510. https://doi.org/10.1086/228311

Gray, M. L., & Suri, S. (2019). *Ghost work: How to stop Silicon Valley from building a new global underclass*. Houghton Mifflin Harcourt.

Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., et al. (2023). How do social media feed algorithms affect attitudes and behavior in an election campaign? *Science, 381*(6656), 398-404. https://doi.org/10.1126/science.abp9364

Habermas, J. (1989). *The structural transformation of the public sphere: An inquiry into a category of bourgeois society* (T. Burger, Trans.). MIT Press. (Original work published 1962)

Habermas, J. (2022). Reflections and hypotheses on a further structural transformation of the political public sphere. *Theory, Culture & Society, 39*(4), 145-171. https://doi.org/10.1177/02632764221112341

Hagiu, A., & Wright, J. (2015). Marketplace or reseller? *Management Science, 61*(1), 184-203. https://doi.org/10.1287/mnsc.2014.2042

Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. *First Monday, 7*(4). https://doi.org/10.5210/fm.v7i4.942

Hoffman, R., & Yeh, C. (2018). *Blitzscaling: The lightning-fast path to building massively valuable companies*. Currency.

Humlum, A., & Vestergaard, E. (2025). *Still waters, rapid currents: Early labor market transformation under generative AI* (NBER Working Paper No. 33777). National Bureau of Economic Research. https://doi.org/10.3386/w33777

International Energy Agency. (2025). Energy and AI. IEA. https://www.iea.org/reports/energy-and-ai

International Energy Agency. (2026). Key questions on energy and AI. IEA. https://www.iea.org/reports/key-questions-on-energy-and-ai

Jack, W., & Suri, T. (2014). Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. *American Economic Review, 104*(1), 183-223. https://doi.org/10.1257/aer.104.1.183

Johnson, G. A., Shriver, S. K., & Goldberg, S. G. (2023). Privacy and market concentration: Intended and unintended consequences of the GDPR. *Management Science, 69*(10), 5695-5721. https://doi.org/10.1287/mnsc.2023.4709

Jones, C. I., & Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review, 110*(9), 2819-2858. https://doi.org/10.1257/aer.20191330

Kamepalli, S. K., Rajan, R., & Zingales, L. (2020). *Kill zone* (NBER Working Paper No. 27146). National Bureau of Economic Research. https://doi.org/10.3386/w27146

Kaplan, S. N., & Strömberg, P. (2003). Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. *The Review of Economic Studies, 70*(2), 281-315. https://doi.org/10.1111/1467-937X.00245

Kattel, R., & Mergel, I. (2019). Estonia's digital transformation: Mission mystique and the hiding hand. In M. E. Compton & P. 't Hart (Eds.), *Great policy successes* (pp. 143-160). Oxford University Press. https://doi.org/10.1093/oso/9780198843719.003.0008

Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *The American Economic Review, 75*(3), 424-440.

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals, 14*(1), 366-410. https://doi.org/10.5465/annals.2018.0174

Kenney, M., & Zysman, J. (2016). The rise of the platform economy. *Issues in Science and Technology, 32*(3), 61-69.

Kerr, W. R., Nanda, R., & Rhodes-Kropf, M. (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives, 28*(3), 25-48. https://doi.org/10.1257/jep.28.3.25

Khan, L. M. (2017). Amazon's antitrust paradox. *The Yale Law Journal, 126*(3), 710-805.

Kirzner, I. M. (1973). *Competition and entrepreneurship*. University of Chicago Press.

Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the age of algorithms. *Journal of Legal Analysis, 10*, 113-174. https://doi.org/10.1093/jla/laz001

Klemperer, P. (1987). Markets with consumer switching costs. *The Quarterly Journal of Economics, 102*(2), 375-394. https://doi.org/10.2307/1885068

Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.

Kostka, G. (2019). China's social credit systems and public opinion: Explaining high levels of approval. *New Media & Society, 21*(7), 1565-1593. https://doi.org/10.1177/1461444819826402

Landes, W. M., & Posner, R. A. (1989). An economic analysis of copyright law. *The Journal of Legal Studies, 18*(2), 325-363. https://doi.org/10.1086/468150

Lehdonvirta, V., & Castronova, E. (2014). *Virtual economies: Design and analysis*. MIT Press. https://doi.org/10.7551/mitpress/9525.001.0001

Lerner, J., & Tirole, J. (2002). Some simple economics of open source. *The Journal of Industrial Economics, 50*(2), 197-234. https://doi.org/10.1111/1467-6451.00174

Lessig, L. (1999). *Code and other laws of cyberspace*. Basic Books.

Lessig, L. (2004). *Free culture: How big media uses technology and the law to lock down culture and control creativity*. Penguin Press.

Lewis, R. A., & Rao, J. M. (2015). The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics, 130*(4), 1941-1973. https://doi.org/10.1093/qje/qjv023

Ley 12/2021, de 28 de septiembre, por la que se modifica el texto refundido de la Ley del Estatuto de los Trabajadores, para garantizar los derechos laborales de las personas dedicadas al reparto en el ámbito de plataformas digitales. Boletín Oficial del Estado, 233, 29 September 2021.

Liebowitz, S. J., & Margolis, S. E. (1994). Network externality: An uncommon tragedy. *Journal of Economic Perspectives, 8*(2), 133-150. https://doi.org/10.1257/jep.8.2.133

Luccioni, A. S., Jernite, Y., & Strubell, E. (2024). Power hungry processing: Watts driving the cost of AI deployment? In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM. https://doi.org/10.1145/3630106.3658542

Marwick, A. E. (2013). *Status update: Celebrity, publicity, and branding in the social media age*. Yale University Press.

Mau, S. (2019). *The metric society: On the quantification of the social*. Polity Press.

Mayer-Schönberger, V., & Cukier, K. (2013). *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt.

Mazzucato, M. (2013). *The entrepreneurial state: Debunking public vs. private sector myths*. Anthem Press.

McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw-Hill.

Nambisan, S. (2017). Digital entrepreneurship: Toward a digital technology perspective of entrepreneurship. *Entrepreneurship Theory and Practice, 41*(6), 1029-1055. https://doi.org/10.1111/etap.12254

Napoli, P. M. (2019). *Social media and the public interest: Media regulation in the disinformation age*. Columbia University Press. https://doi.org/10.7312/napo18454

Nedelkoska, L., & Quintini, G. (2018). *Automation, skills use and training* (OECD Social, Employment and Migration Working Papers No. 202). OECD Publishing. https://doi.org/10.1787/2e2f4eea-en

Negroponte, N. (1995). *Being digital*. Alfred A. Knopf.

Nieborg, D. B., & Poell, T. (2018). The platformization of cultural production: Theorizing the contingent cultural commodity. *New Media & Society, 20*(11), 4275-4292. https://doi.org/10.1177/1461444818769694

Nielsen, R. K., & Ganter, S. A. (2018). Dealing with digital intermediaries: A case study of the relations between publishers and platforms. *New Media & Society, 20*(4), 1600-1617. https://doi.org/10.1177/1461444817701318

Norris, P. (2001). *Digital divide: Civic engagement, information poverty, and the Internet worldwide*. Cambridge University Press. https://doi.org/10.1017/CBO9781139164887

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187-192. https://doi.org/10.1126/science.adh2586

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.

Oberholzer-Gee, F., & Strumpf, K. (2007). The effect of file sharing on record sales: An empirical analysis. *Journal of Political Economy, 115*(1), 1-42. https://doi.org/10.1086/511995

OECD. (2020). *A roadmap toward a common framework for measuring the digital economy*. OECD Publishing. https://www.oecd.org/digital/ieconomy/roadmap-toward-a-common-framework-for-measuring-the-digital-economy.pdf

OECD. (2026, January). AI use by individuals surges across the OECD as adoption by firms continues to expand. OECD. https://www.oecd.org/en/about/news/announcements/2026/01/ai-use-by-individuals-surges-across-the-oecd-as-adoption-by-firms-continues-to-expand.html

Osterwalder, A., & Pigneur, Y. (2010). *Business model generation: A handbook for visionaries, game changers, and challengers*. John Wiley & Sons.

Oviatt, B. M., & McDougall, P. P. (1994). Toward a theory of international new ventures. *Journal of International Business Studies, 25*(1), 45-64. https://doi.org/10.1057/palgrave.jibs.8490193

Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin Press.

Parker, G. G., & Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management Science, 51*(10), 1494-1504. https://doi.org/10.1287/mnsc.1050.0400

Parker, G. G., Van Alstyne, M. W., & Choudary, S. P. (2016). *Platform revolution: How networked markets are transforming the economy and how to make them work for you*. W. W. Norton.

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press. https://doi.org/10.4159/harvard.9780674736061

Perez, C. (2002). *Technological revolutions and financial capital: The dynamics of bubbles and golden ages*. Edward Elgar. https://doi.org/10.4337/9781781005323

Pickard, V. (2020). *Democracy without journalism? Confronting the misinformation society*. Oxford University Press. https://doi.org/10.1093/oso/9780190946753.001.0001

Pigou, A. C. (1920). *The economics of welfare*. Macmillan.

Poell, T., Nieborg, D. B., & Duffy, B. E. (2022). *Platforms and cultural production*. Polity Press.

Polski Standard Płatności. (2026). *Blisko 3 mld transakcji BLIKIEM w 2025 r. i ponad 2 mln nowych użytkowników* [Nearly 3 billion BLIK transactions in 2025 and over 2 million new users] [Press release]. https://www.blik.com/blisko-3-mld-transakcji-blikiem-w-2025-r-i-ponad-2-mln-nowych-uzytkownikow

Prawo.pl. (2022). *Allegro zapłaci 210 mln zł kary za faworyzowanie swojego sklepu i zakazane klauzule* [Allegro to pay PLN 210 million for favouring its own store and prohibited clauses]. https://www.prawo.pl/prawo/allegro-210-mln-zl-kary-za-faworyzowanie-swojego-sklepu,519064.html

Rahman, H. A. (2021). The invisible cage: Workers' reactivity to opaque algorithmic evaluations. *Administrative Science Quarterly, 66*(4), 945-988. https://doi.org/10.1177/00018392211010118

Resnick, P., & Zeckhauser, R. (2002). Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. In M. R. Baye (Ed.), *The economics of the Internet and e-commerce* (Advances in Applied Microeconomics, Vol. 11, pp. 127-157). Emerald. https://doi.org/10.1016/S0278-0984(02)11030-3

Reuters Institute for the Study of Journalism. (2026). Digital News Report 2026. University of Oxford. https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2026

Ries, E. (2011). *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. Crown Business.

Roberts, S. T. (2019). *Behind the screen: Content moderation in the shadows of social media*. Yale University Press.

Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association, 1*(4), 990-1029. https://doi.org/10.1162/154247603322493212

Rohlfs, J. (1974). A theory of interdependent demand for a communications service. *The Bell Journal of Economics and Management Science, 5*(1), 16-37. https://doi.org/10.2307/3003090

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy, 98*(5, Part 2), S71-S102. https://doi.org/10.1086/261725

Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication, 10*, 3758-3784.

Sadowski, J. (2019). When data is capital: Datafication, accumulation, and extraction. *Big Data & Society, 6*(1), 1-12. https://doi.org/10.1177/2053951718820549

Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics, 36*(4), 387-389. https://doi.org/10.2307/1925895

Saxenian, A. (1994). *Regional advantage: Culture and competition in Silicon Valley and Route 128*. Harvard University Press.

Scholz, T. (2016). *Platform cooperativism: Challenging the corporate sharing economy*. Rosa Luxemburg Stiftung, New York Office.

Schumpeter, J. A. (1934). *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle*. Harvard University Press.

Schumpeter, J. A. (1942). *Capitalism, socialism and democracy*. Harper & Brothers.

Shane, S., & Venkataraman, S. (2000). The promise of entrepreneurship as a field of research. *Academy of Management Review, 25*(1), 217-226. https://doi.org/10.5465/amr.2000.2791611

Shapiro, C., & Varian, H. R. (1999). *Information rules: A strategic guide to the network economy*. Harvard Business School Press.

Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, communications, and the public interest* (pp. 37-72). Johns Hopkins Press.

Smythe, D. W. (1977). Communications: Blindspot of Western Marxism. *Canadian Journal of Political and Social Theory, 1*(3), 1-27.

Solow, R. M. (1987, July 12). We'd better watch out [Review of the book *Manufacturing matters*, by S. S. Cohen & J. Zysman]. *The New York Times Book Review*, 36.

Spigel, B. (2017). The relational organization of entrepreneurial ecosystems. *Entrepreneurship Theory and Practice, 41*(1), 49-72. https://doi.org/10.1111/etap.12167

Srnicek, N. (2017). *Platform capitalism*. Polity Press.

Stam, E. (2015). Entrepreneurial ecosystems and regional policy: A sympathetic critique. *European Planning Studies, 23*(9), 1759-1769. https://doi.org/10.1080/09654313.2015.1061484

Stanford Institute for Human-Centered Artificial Intelligence. (2026). The 2026 AI Index report. Stanford University. https://hai.stanford.edu/ai-index/2026-ai-index-report

Sundararajan, A. (2016). *The sharing economy: The end of employment and the rise of crowd-based capitalism*. MIT Press.

Sussan, F., & Acs, Z. J. (2017). The digital entrepreneurial ecosystem. *Small Business Economics, 49*(1), 55-73. https://doi.org/10.1007/s11187-017-9867-5

Tadelis, S. (2016). Reputation and feedback systems in online platform markets. *Annual Review of Economics, 8*, 321-340. https://doi.org/10.1146/annurev-economics-080315-015325

Tapscott, D. (1996). *The digital economy: Promise and peril in the age of networked intelligence*. McGraw-Hill.

Teece, D. J. (2010). Business models, business strategy and innovation. *Long Range Planning, 43*(2-3), 172-194. https://doi.org/10.1016/j.lrp.2009.07.003

Terranova, T. (2000). Free labor: Producing culture for the digital economy. *Social Text, 18*(2), 33-58. https://doi.org/10.1215/01642472-18-2_63-33

Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc., No. 1:20-cv-00613-SB (D. Del. Feb. 11, 2025).

Uber BV v Aslam [2021] UKSC 5.

UNCTAD. (2019). *Digital economy report 2019: Value creation and capture - implications for developing countries*. United Nations. https://doi.org/10.18356/c7dc937a-en

United States v. Microsoft Corp., 253 F.3d 34 (D.C. Cir. 2001).

Vallas, S., & Schor, J. B. (2020). What do platforms do? Understanding the gig economy. *Annual Review of Sociology, 46*, 273-294. https://doi.org/10.1146/annurev-soc-121919-054857

van Dijck, J., Poell, T., & de Waal, M. (2018). *The platform society: Public values in a connective world*. Oxford University Press. https://doi.org/10.1093/oso/9780190889760.001.0001

van Dijk, J. A. G. M. (2020). *The digital divide*. Polity Press.

Van Parijs, P., & Vanderborght, Y. (2017). *Basic income: A radical proposal for a free society and a sane economy*. Harvard University Press. https://doi.org/10.4159/9780674978072

Varian, H. R. (1989). Price discrimination. In R. Schmalensee & R. D. Willig (Eds.), *Handbook of industrial organization* (Vol. 1, pp. 597-654). North-Holland. https://doi.org/10.1016/S1573-448X(89)01013-7

Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives, 28*(2), 3-28. https://doi.org/10.1257/jep.28.2.3

von Briel, F., Davidsson, P., & Recker, J. (2018). Digital technologies as external enablers of new venture creation in the IT hardware sector. *Entrepreneurship Theory and Practice, 42*(1), 47-69. https://doi.org/10.1177/1042258717732779

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science, 359*(6380), 1146-1151. https://doi.org/10.1126/science.aap9559

Waldfogel, J. (2017). How digitization has created a golden age of music, movies, books, and television. *Journal of Economic Perspectives, 31*(3), 195-214. https://doi.org/10.1257/jep.31.3.195

Wardle, C., & Derakhshan, H. (2017). *Information disorder: Toward an interdisciplinary framework for research and policymaking* (Council of Europe Report DGI(2017)09). Council of Europe.

Webster, J. G. (2014). *The marketplace of attention: How audiences take shape in a digital age*. MIT Press. https://doi.org/10.7551/mitpress/9892.001.0001

Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society, 33*(1), 56-75. https://doi.org/10.1177/0950017018785616

Wu, T. (2003). Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law, 2*, 141-179.

Wu, T. (2016). *The attention merchants: The epic scramble to get inside our heads*. Alfred A. Knopf.

Wu, T. (2018). *The curse of bigness: Antitrust in the new gilded age*. Columbia Global Reports. https://doi.org/10.2307/j.ctv1fx4h9c

Zhu, F., & Liu, Q. (2018). Competing with complementors: An empirical look at Amazon.com. *Strategic Management Journal, 39*(10), 2618-2642. https://doi.org/10.1002/smj.2932

Zott, C., Amit, R., & Massa, L. (2011). The business model: Recent developments and future research. *Journal of Management, 37*(4), 1019-1042. https://doi.org/10.1177/0149206311406265

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.

Zucker, L. G. (1986). Production of trust: Institutional sources of economic structure, 1840-1920. *Research in Organizational Behavior, 8*, 53-111.

ABOUT THE AUTHORS

Henrietta Nagy PHD, HABIL.

Henrietta Nagy, PhD, Habil. is a Professor of Economics and Management at Milton Friedman University in Budapest, Hungary. With more than two decades of experience in higher education, her research focuses on regional policy and planning, agricultural economics, and rural sustainable development. She is actively involved in international academic cooperation and serves on scientific advisory bodies across Central Europe.

Tamás Tóth PHD, HABIL.

Tamás Tóth, PhD, Habil. is a Professor of Economics and Management and Head of the Department of Economics and Management Sciences at Milton Friedman University. Throughout his academic career, he has served in various senior university leadership roles, including Deputy Dean, Vice-Rector, and head of doctoral programs. His primary research interests encompass management and organizational sciences, regional economics, and higher education development.

Jarosław Kinal PHD

Jarosław Kinal, PhD is a sociologist and academic researcher at the University of Rzeszów, Poland. He has extensive experience in academic administration and international academic cooperation. His research interests include contemporary social structures, sociology of communication, cultural change, and regional studies.

ISBN 978-963-9559-80-6