Contents

Dedication ix

Foreword by Richard Susskind xi

A personal note xiii
The philosopher and the practitioner xiv
Thinking in trajectories xv
What I believe about AI xvi
The river and the riverbed xvi

About the author xix

How to read this book xxi
The architecture xxi
The four parts xxi
The canvases xxii
Paths through the book xxiii
A note on the composite narrative xxiv
A note on the real-world events xxiv
One last thing xxv

Introduction xxvii
The new normal xxviii
The system is your superpower xxix
What this means for in-house legal teams xxx
What this means for law firms xxxi
What this means for clients xxxii
Who this book is for xxxii
The strategic task ahead xxxiii

Part One: The Case for Why 1
The Master Principle: The Why Equation 1
The Why Equation canvas and index 2
A counterintuitive insight 2

Chapter 1: Purpose – from shiny objects to meaningful decisions 3
A Tuesday morning meeting 3
Why this keeps happening 4
The Principle: Purpose 7
The method – finding your purpose 13
Case study: Axiom's exit from the loop 15
The human factor – the Anxiety Tax 16
Your canvas so far 16
What's next? 17
Key takeaways 18

Chapter 2: Clarity – from foggy mandates to measurable commitments 21
The Friday email 21
hy this keeps happening 23
The principle: clarity 25
The method – building shared clarity 30
Case study: Harvey's self-improving auto-research agent 32
The human factor – the invisible labor 33
Your canvas so far 33
The Why Equation Canvas 34
What's next? 34
Key takeaways 35

Chapter 3: Alignment – from innovation theater to active championship 39
A third attempt 39
Why this keeps happening 40
The Principle: Alignment 43
The method – building a coalition 48
Case study: Shoosmiths' million-prompt bet 49
The human factor – what the change actually costs 50
The Why Equation Canvas 51
Your Readiness Index 53

Graystone's three pilot failures 54
End of Part 1: The Why Equation complete 56
Key takeaways 57

Part Two: Foundations First 61
The Master Principle: The Integration Function 61
Why the order matters 61
The foundation canvas and index 62
A counterintuitive insight 62

Chapter 4: Infrastructure as strategy – from cost center to competitive architecture 63
Six months, nine days 63
Why infrastructure is the first foundation 63
Why legacy infrastructure made sense 63
How AI decenters human operators 64
What legal AI infrastructure should look like 65
The portability imperative 71
The Foundation Canvas 72
Your canvas so far 74
What's next? 74
Key takeaways 75

Chapter 5: Knowledge as product – from tacit expertise to explicit knowledge 79
The hundredth question 79
Why knowledge is the second foundation 79
Why legacy knowledge management made sense 80
How AI upends tacit knowledge 80
From tacit knowledge to compounding asset 81
The compounding advantage in practice 86
The Foundation Canvas 86
Your canvas so far 89
What's next? 89
Key takeaways 90

Chapter 6: Guardrails as accelerator – from compliance burden to innovation engine 93
The call from Meridian 93
Why guardrails complete the trilogy 94

Why legacy risk management made sense 94
How AI disrupts security controls 94
What trustworthy AI infrastructure looks like 96
The compounding effect of trust 100
The Foundation Canvas 101
The running example – Panel 3: Defensibility (Graystone Example) 102
Your Foundation Index 104
Graystone's three foundation failures 105
End of Part Two: The Foundation complete 107
Key takeaways 108

Part Three: Capability Development 113
The Master Principle: The Transformation Triangle 113
The Transformation Triangle canvas and index 114
A counterintuitive insight 114

Chapter 7: Reactive to proactive – from intellectual labor to expert design 115
The Sunday night text 115
The reactive legal service delivery model 115
Intellectual labor and expert design 117
The Service Transformation Principle 118
The three levers of transformation 120
Case study: Macfarlanes Amplify 124
Non-billable time and visible wins 125
The Transformation Triangle Canvas 126
What's next? 129
Key takeaways 129

Chapter 8: Bespoke to flywheel – from episodic deliverables to compounding capital 133
The folder no one will open 133
Why bespoke persists 133
When labor became cheap and capital cheaper 134
The flywheel principle 135
The legal flywheel has six steps: 135
Case Study: LegalQuants 137
Managed services – how capital earns 139
The lock-in question 140

The Transformation Triangle Canvas 141
Your Canvas so far 144
What's next? 145
Key takeaways 146

Chapter 9: Training to learning – from projecting expertise to feedback loops 149
Two hundred and thirty-seven 149
Why training has always failed 150
Training and learning are different things 150
The Feedback Engine 152
The loops that touch the world 153
Loop One: The Client Signal Loop 153
Loop Two: The Practice Signal Loop 154
Loop Three: The Outcome Signal Loop 154
A note on apprenticeship 155
Case study: Freshfields and the architecture of learning 156
The Transformation Triangle Canvas 158
Your Capability Index 161
Graystone's three triangle failures 162
A firm's learning rate is its competitive moat 164
End of Part Three: The Transformation Triangle complete 164
Key takeaways 165

Part Four: Identity, Economics, and Discipline 167
The Master Principle: The Compounding Loop 167
The four phases of transformation 168

Chapter 10: The Why-first law firm – the new identity 169
Three firms 169
The category error – AI-first and AI-native law firms 172
What the firm becomes – an antifragile legal platform 174
The five disruptions 175
The decade ahead 177

Chapter 11: The Why-first business model – the new economics 183
One page 183
The Value Principle 185
The architecture 187
What changes for the lawyer 191

What changes for the buyer 192

Chapter 12: The Why-first blueprint – the new discipline 197
A beginning 197
The Transformation Principle 198
The Why-first Blueprint 198
Phase 1: Direction 199
Phase 2: Foundation 200
Phase 3: Transformation 202
Phase 4: Identity 203
Timing and pace 205
Monday morning 205
The recursive instrument 206
The river 208

Appendix: Frameworks and principles 211
The architecture at a glance 211
Part One: The Case for Why 212
Part Two: Foundations First 213
Part Three: Capability Development 214
Part Four: Identity, Economics, and Discipline 215
The Principles across the book 217
The Canvases 217

Further reading 219

Acknowledgements 223

About Globe Law and Business 227

Dedication

For Miriam, Julia and Maya, my loves, and my reason why.

Foreword

Richard Susskind

No one knows where AI is taking us, or perhaps where we are taking AI. We do know, however, that the underpinning technologies – from the chips to the neural networks – are advancing at a ferocious pace. And we know too that much of the legal world is in a frenzy about the impact of AI.

Most of the current focus on legal AI is on the next couple of years. Lawyers, suppliers, and clients are largely preoccupied today with the efficiency and productivity gains that the latest systems might deliver. At the same time, an army of commentators are disgorging an endless stream of upbeat blogs and posts, many of which are of course ghost-written by AI. The sceptics have a strong voice too, arguing from their hearts that many of the predictions being made are overblown and of considerable concern.

Into this frenzy of argument and activity comes the good sense, clear thinking, and authentic voice of Tom Martin, who writes with the heart of a philosopher, the mind of a lawyer, and the instincts of a strategist. *AI with Purpose* encourages and requires its readers to think more deeply about how AI can be better deployed in the legal world. It rises above the press releases and the pilots to provide a systematic guide to legal AI.

This book should help legal professionals and their leaders to think more clearly about the impact of AI on legal services and to act more confidently. It is at once practical and theoretical – there are countless tips for the action-oriented, but the author also provides frameworks and models that should inject greater rigor into the adoption of AI in law.

Tom Martin's ideas are brought to life in his fictional discussion of Graystone LLP and in his non-fictional case studies of high-profile projects, systems, and businesses. He provides us with principles and paradoxes that stretch and shape our understanding. He writes with enthusiasm and urgency. His passion for the field is evident on every page.

We should remind ourselves of what is at stake. The law is our most important method of social control. Our current system – the legal profession, law schools, courts – is creaking. An upgrade is long overdue. In legal AI, many of us can see a key to great improvement. But we are still at the foothills, feeling

our way to better places. Our efforts to transform the law should be made responsibly and methodically. In this context, Tom Martin's work is an important contribution.

I wish this book every success.

Professor Richard Susskind CBE KC (Hon)
President, Society for Computers and Law
Author, *The Future of Law*, *Tomorrow's Lawyers*, and *How to Think About AI*

A personal note

"No man ever steps in the same river twice, for it is not the same river and he is not the same man."
Heraclitus, Fragment 91

I studied continental philosophy as an undergraduate at Yale University. Not as preparation for law school, though that is what it eventually became. I studied philosophy because I wanted to understand the larger questions of existence, like why we are here, how things change, and what, if anything, persists through transformation.

The philosopher who made the greatest impression on me was Heraclitus of Ephesus, who wrote 25 centuries ago and all we have left of his writings now are fragments. His central insight was that change is the only constant. You cannot step in the same river twice. Not because the river has moved on, but because you have too. The water is different. The current has shifted. And the person standing on the bank is not the same person who stood there a moment ago, because the act of stepping into the river has already changed them.

I think about Heraclitus often when I think about AI and the law and this particular moment in our history. Every conversation I have with legal professionals about artificial intelligence takes place in a river that is moving. The models have improved since last week. A new tool has launched since last month. A regulatory opinion has shifted since the last conference. The technology you evaluated in January is not the technology available in March. And you are not the same professional you were when you started evaluating it, because the act of engaging with AI has already begun to change how you think and how you work.

This is perplexing.

But, the disorientation is not a bug of the current moment. It is the new reality of the current moment. Rather than fight against the current, we must learn to navigate it with purpose.

That is what this book is for.

The philosopher and the practitioner

Before I explain what this book argues, I want to acknowledge an intellectual debt.

I first met Professor Richard Susskind in 2018, at the British Legal Technology Forum, where he moderated a panel discussion of which I was a speaker, on "The Rise of AI Chatbots". Then, in the summer of 2025, I had the privilege of interviewing him for my podcast, LawDroid Manifesto. I called the episode "The AI Philosopher" and I meant it as a compliment.

For four decades, Susskind has been the most important thinker about the future of legal services, and his latest book, *How to Think About AI*, is the clearest framework I have encountered for understanding what artificial intelligence means, not just for lawyers, but for the professions, for society, and for what it means to be an intelligent *human* being in an age of reasoning machines.

Susskind draws a distinction that runs beneath everything in the book you are holding. He distinguishes between process-thinking and outcome-thinking. Process-thinkers ask how a system works. Outcome-thinkers ask what it produces. In AI, the process-thinkers point out, correctly, that large language models do not reason the way humans reason – they predict tokens, not truth. The outcome-thinkers point out, also correctly, that the outputs are increasingly indistinguishable from expert human work. Susskind does not resolve this tension. He insists we hold both perspectives simultaneously, because the tension itself is where wisdom lives.

He also distinguishes three ways AI changes professional work – automation, innovation, and elimination. Automation grafts technology onto existing processes, making them faster and cheaper. Innovation uses technology to deliver outcomes in ways that were previously impossible – online dispute resolution rather than physical courts, embedded compliance rather than periodic audits. Elimination removes the need for the professional service altogether, the way the automobile eliminated the Great Manure Crisis of the 1890s – not by cleaning up the horse dung more efficiently, but by making horses irrelevant to transportation.

Susskind's framework is the *philosophical* architecture for this book. He identifies the destination – a world where legal services are not merely automated but reimagined and, in some domains, rendered unnecessary by upstream prevention. What his framework does not provide is the *operational* architecture for getting there. How does a managing partner translate "innovation, not just automation" into a budget request? How does a legal

operations leader build the infrastructure that makes elimination possible? How does a practice group leader convince skeptical partners that the goal is not to work faster but to redesign what work means?

Those are the questions this book answers. If Susskind describes the river's destination, I am trying to teach you how to swim.

Thinking in trajectories

Heraclitus's river carries one more lesson for this book.

Most people evaluate AI based on what it can do *today*. They test the current model. They assess the current tool. They judge the current output against the current standard of professional work. And they conclude, often reasonably (for today), that AI is not yet ready, that it hallucinates too often, misses too much nuance, requires too much supervision to be worth the trouble.

This evaluation is accurate and irrelevant. It is accurate because current AI systems do have real limitations. It is irrelevant because you are not deciding whether to adopt today's AI. You are deciding whether to build the organizational capacity to adopt AI as it improves, which it will, rapidly, and in ways that will make today's limitations seem quaint.

Susskind makes this point forcefully. In the 1980s, he wrote his doctorate on AI and law at Oxford. The systems he built then were painstaking decision trees that required years of expert knowledge engineering to encode a fraction of what a junior lawyer knows intuitively. He has watched AI evolve from those brittle beginnings through decades of incremental progress to the current moment, where a general-purpose language model can pass the bar exam and draft competent legal analysis on the first attempt. His conclusion: we are at the foothills. The current systems, impressive as they are, are not the technology that will transform the profession. They are the early drafts of that technology.

As Susskind puts it, we must ask ourselves, "What if AGI?".[1] Because when we ask ourselves, "If Artificial General Intelligence already existed, how would we approach things differently?", we free ourselves from the blinders of only seeing what AI is capable of now.

Thinking in trajectories rather than snapshots changes the strategic calculus entirely. If you evaluate AI as a snapshot of this moment in its evolution, you make a tool decision – is this product worth buying today? If you evaluate AI as a trajectory, you make an infrastructure decision – are we building the organizational capacity to adopt whatever comes next, faster than our competitors?

That is an infrastructure question, a knowledge architecture question, a governance question, a cultural question. It is, in other words, every question this book addresses.

What I believe about AI

I believe that AI is the most important technology to arrive in the legal profession in my lifetime, and possibly in the profession's history. I believe it is changing not just how legal work is performed but what legal services are, how they are delivered, who can access them, and what the profession exists to do. I believe this change is inevitable but not automatic; it has become a force multiplier for organizations that continually prepare for it, and an extinction level event for those who do not.

I also believe that technology alone changes nothing. I have watched brilliant AI tools fail because no one defined what success meant. I have watched mediocre tools succeed because they were attached to clear workflows and accountable owners. I have watched organizations spend millions on platforms they never adopted and thousands on simple automations that transformed how they worked. The pattern is consistent – the technology is the easy part. The hard part is everything around it.

And I believe (this is the philosophy major talking) that the most important question in any technology adoption is not "What can it do?" but "Why?" Purpose is not a strategic nicety. It is an ontological requirement. A tool without purpose is not a tool. It is a distraction.

This book is built on these beliefs. It is prescriptive, not descriptive. It tells you what to do, in what order, and why. It provides frameworks, diagnostic patterns, principles, and practical guidance for every stage of AI transformation from initial purpose-setting through market leadership.

I do not pretend that transformation is easy, or fast, or guaranteed. I do not gloss over the human costs of change – the anxiety, the identity disruption, the loss of purpose. I have set out the human impacts that transformation has on the people who undertake it, because I have learned that the organizations that forget about their people do not succeed.

The river and the riverbed

One last thought from Heraclitus.

The fragment about the river is usually interpreted as a statement about change – everything flows, nothing remains the same. But there is another way to read it. The river changes constantly, the water is always new, the

current always shifting. But the riverbed changes slowly. The underlying geology, the gradient, the fundamental shape of the channel – these persist across seasons and years, even as the surface is in constant motion.

AI is the river. The models change. The tools change. The capabilities expand. What you can do with AI today is different from what you could do six months ago and different from what you will be able to do six months from now. If you try to build your strategy on the surface of the river, on specific tools, specific models, specific capabilities, you will be rebuilding constantly, chasing a current that never holds still.

But the riverbed – the organizational foundations of purpose, clarity, alignment, infrastructure, knowledge architecture, governance, service design, and learning systems – is foundational. These are the structures that persist regardless of which AI models power them. Build on the riverbed, and every new wave of capability flows through the channel you have already cut. Build on the surface, and you will be swept away by the current – reactive, exhausted, and perpetually starting over.

This book teaches you to build on the riverbed.

The rushing waters will come and go, but your convictions, principles, judgment and humanity will be the rockbed that stands the test of time.

Tom Martin
Vancouver, British Columbia, Canada
July 2026

Reference

1 Keynote – What if AGI? (Prof. Richard Susskind). Thinking About Thinking. 3 March 2026. www.youtube.com/watch?v=c5YHiNsR5BY. Talk was given at the Algorithmic Innovation and Entrepreneurship Global Summit on Open Problems for AI (London, 2025).

About the author

Thomas G. Martin is at the forefront of transforming the legal industry. As CEO and founder of LawDroid, a pioneering legal AI company, he defines the AI technological frontier. As principal of Deep Legal, a legal AI transformation consultancy, he advises enterprises on strategy, AI adoption, and operating model redesign. As co-founder of the Academy of Legal Innovation, which yearly hosts the American Legal Technology Awards, he honors legal innovators across North America. Finally, as an adjunct professor at Suffolk Law School, he educates law students about Generative AI and the delivery of legal services.

He has been recognized as an ABA Legal Rebel and Fastcase 50 Honoree, appointed by the President of the American Bar Association to the Center for Innovation Governing Council.

Tom is "The AI Law Professor" for a regular column published by the Thomson Reuters Institute, as well as contributing to publications such as the *National Law Review*. Tom has also captivated audiences at renowned events like ABA TechShow, FutureLaw (Estonia), Global LegalTech Hub (Spain), LegalWeek, ILTACON, ClioCon, and Legal Innovators. Tom hosts the LawDroid Manifesto newsletter and podcast with 6,500+ subscribers and the annual LawDroid AI Conference with 1,500+ global attendees.

He serves as a mentor at the Yale Tsai Center for Innovative Thinking and ATJ Tech Fellows. A graduate of Yale University and UCLA School of Law, Tom is a proud father of two amazing daughters and lives in Vancouver, Canada.

How to read this book

The architecture

This book is designed to be read in order. The 12 chapters form a cumulative guide in which each builds on what came before. The Why Equation in Part One opens the gates of strategy. The Foundation Canvas in Part Two builds the base that strategy requires to execute. The Transformation Triangle in Part Three runs the operations the foundation supports. And the Why-First Blueprint in Part Four sequences direction, foundation, transformation, and identity into a single phased program that a managing partner can run on Monday morning. If you skip ahead, you'll learn how it all comes together, but you will miss out on the why and how.

That said, I know how busy people read professional books – selectively, under time pressure, looking for the chapter that solves Tuesday's problem today. So here is a guide to the structure that will help you navigate efficiently without losing the thread.

The four parts

Part one: The case for why (chapters 1-3): Purpose, Clarity, Alignment

These three chapters answer the questions every AI initiative must address before a single tool is selected: Why are we doing this, What exactly will we accomplish, and Is everyone genuinely committed to making it work? Together they form the *Why Equation*: Purpose plus Clarity plus Alignment equals Transformation Readiness. If your organization is struggling with AI initiatives that launch but do not land, start here. These chapters diagnose the root cause.

Part two: Foundations first (chapters 4-6): Infrastructure, Knowledge, Guardrails

These chapters address what must exist before AI capability can function at the speed the market is now producing it – the infrastructure stack that makes deployment fast, the structured expertise that makes deployment

useful, and the governance architecture that makes deployment defensible. They are governed by the *Integration Function* – a firm absorbs new capability at the rate its existing conditions permit, and no faster. If your organization has purpose and alignment but still cannot get AI tools into production, these chapters explain why.

Part three: Capability development (chapters 7-9): Service transformation, Flywheels, Learning systems
These chapters describe the new operating model that AI makes possible – proactive service delivery rather than reactive response, software that compounds across engagements rather than work that resets to zero, and learning systems that absorb signal from the contact surface (the firm's live matters, its lawyers' day-to-day use of its tools, and the outcomes that follow) rather than training programs that project outward. Together they form the *Transformation Triangle* of Expertise, Tools, and Learning. If your organization has built the infrastructure but has not yet changed how work is done or how people develop, these chapters show what comes next.

Part four: Identity, economics, and discipline (chapters 10-12): The Why-first firm, The business model, The blueprint
These chapters address what holds the components together – an identity that absorbs new capability rather than mutating with it, an economic model that prices what the firm has become rather than how long it spent becoming it, and a phased program that sequences the work in the order it requires. The four phases of the program are *Direction, Foundation, Transformation, and Identity*, mapping directly onto the four parts of this book. If you want to see the full journey from start to finish, with deliverables, exit criteria, and failure modes for each phase, Part Four is the destination.

The canvases

The canvases are the operational deliverable of the book. The chapters are how you learn to draft them.

Each canvas is an artifact you build:
The *Why Equation Canvas* in part one.
The *Foundation Canvas* in part two.
The *Transformation Triangle Canvas* in part three.
The *Why-first Blueprint* in part four.

Each canvas has three fields. Each field is drafted in one chapter, against a running example you follow across the whole part. By the time you finish a part, the Canvas is complete, the example is intact, and you have an artifact you can use in your firm against your own initiative.

Each canvas (in parts one through three) also has an index: the *Readiness Index* at the end of part one, the *Foundation Index* at the end of part two, and the *Capability Index* across part three. The index is the readout. It reads the *minimum* of the three field scores, not the average, because each part of the book argues, in its own register, that the relevant integration is multiplicative rather than additive. A firm strong on two fields and weak on the third is a weak-field firm, not an average-field firm. The lowest score is the bottleneck.

Each chapter also points to toolkits available on the book's online companion site: aiwithpurposebook.com. The toolkits are the worksheets, scoring rubrics, and worked templates that accompany the chapter's framework. They are not required reading. They are required *doing*, when you are ready to apply the framework to your own organization.

Paths through the book

While I recommend reading in order, here are accelerated paths for readers with specific needs.

- *The managing partner or general counsel who needs to understand why AI initiatives keep disappointing:* Chapters 1, 2, and 3 for the Why Equation, then chapter 10 for the Why-first identity, chapter 11 for the economic transformation that identity demands, and chapter 12 for the phased program that sequences the work.
- *The chief innovation officer or head of legal operations who has been asked to "make AI work":* Chapters 4, 5, and 6 for infrastructure, knowledge, and guardrails, then chapter 9 for the Feedback Engine that turns deployed capability into compounding learning, and chapter 12 for the sequencing discipline that holds the program together over years rather than quarters.
- *The practice group leader who wants to transform how their team delivers work:* Chapters 7 and 8 for the shift from reactive to proactive service delivery and the flywheel model that turns each engagement into an investment in the next. Then read chapter 11 for the pricing models that make transformation economically sustainable, because services redesigned around outcomes cannot be priced by the hour.
- *The law student or early-career lawyer who wants to understand what*

AI means for their career: start with the Human Factor callouts that tell the human story of transformation, Chapter 5 to learn about knowledge as skills acquisition, and Chapter 9 for insight as to what the new apprenticeship model looks like.

- *The legal aid leader who serves people rather than corporations:* every chapter applies to your work, often more urgently than to the firms and departments that dominate the legal technology conversation. The proactive service models of chapter 7, the flywheel systems of chapter 8, and the economic frameworks of chapter 11 describe the future of legal aid delivery as much as they describe the future of corporate law. Where the book discusses competitive advantage, read access to justice. The frameworks are the same; the stakes are higher.

A note on the composite narrative

The firm at the center of this book, Graystone LLP, is fictional. So is Sarah Chen, the general counsel at Meridian Health Systems. So are the associates, partners, IT directors, and innovation officers who populate the chapter openings. They are drawn from real people and real situations, but no character represents a specific individual and no scene depicts a specific event.

I made this choice not to obscure but to clarify. Real organizations have dozens of complicating factors that would distract from the principles each chapter teaches. Composites let me isolate the dynamics that matter. If you recognize your own organization in these pages, that is because the patterns are common enough to be universal, not because I am describing your firm.

The chapters in parts one, two, and three follow Graystone through a sequence of failures and recoveries, each anchored in a different protagonist whose name appears in the running example. This is deliberate. The reader who follows Graystone through 12 chapters watches a single firm pass through every failure mode the book describes, and arrive, by part four, at the firm those failures were teaching it to become.

A note on the real-world events

Unlike the composites, the real-world events described in this book are reported as accurately as I could manage at the time of writing. Harvey's valuation, Legora's funding, Macfarlanes' Amplify, Allen & Overy's ContractMatrix, Shoosmiths' million-prompt bonus, Ropes & Gray's TrAIlblazers program, the OpenClaw phenomenon, the Claude-Native law firm, the Legal Quants community, the vibe coding movement, the launch of Claude Cowork, and

the Freshfields partnership with Anthropic are matters of public record. I have used them to ground the book's frameworks in observable reality.

But the world of AI moves fast. By the time you read this, some of these events will have developed further. Valuations will have changed. New tools will have launched. New controversies will have emerged. That is fine. The events are illustrations, not foundations. The frameworks they illustrate will outlast the specific companies and products that populate them.

One last thing

This book contains a great many frameworks and principles. Martin's Paradoxes. The Why Equation. The Foundation. The Transformation Triangle. The Why-first Law Firm. The Shiny Object Loop. The Fog Cascade. The Innovation Theater Loop. The Bespoke Trap. The Integration Function. The Compounding Loop. The Value Principle. The Transformation Principle. These names exist because strategy that cannot be remembered cannot be applied. For your convenience, I summarize them in *Frameworks and Principles* at the end of the book.

I want you to walk out of a partner meeting and say, *"We have a clarity problem,"* and have everyone in the room know what you mean. I want you to look at a struggling pilot and diagnose it as Innovation Theater. I want you to hear someone say *"We need an AI strategy"* and recognize the opening bars of the Shiny Object Loop. I want you to read your firm's next strategic plan and notice, in the gap between what it claims and what your operations actually do, an alignment gap surfacing.

The frameworks are tools. Use them as tools. Adapt them to your context. Argue with them where they do not fit. But start with them, because starting from a framework is faster than starting from scratch, which, if you think about it, is the entire point of this book.

Let us begin.

Introduction

Thomas Martin

"We're living in a moment as transformational as the drafting of the Constitution."[1]

It was November 2025, just three short years after the launch of ChatGPT, and Richard Tromans had opened the Legal Innovators New York conference with a bold statement. It was the kind of statement that might have sounded hyperbolic a year earlier. It did not anymore. The evidence was already accumulating around us.

New capabilities were arriving faster than legal organizations could absorb them. Lawyers were building working software through conversation. Legal AI companies were attracting multibillion-dollar valuations. The world's largest technology companies were moving directly into legal work. Open-source projects were reproducing features that had recently appeared to require vast engineering teams and venture funding. The legal market had entered a period in which the boundaries between professional expertise, software, institutional knowledge, and organizational design were beginning to dissolve.

For most of its history, the legal profession was treated as a technological backwater. It was a specialized market selling niche tools to a profession, governed by *stare decisis*, that did not particularly want to change. Buyers were cautious, sales cycles were long, budgets were limited, and technology departments were expected to keep systems running rather than shape the future of the organization. Innovation was something law firms discussed at conferences, demonstrated in pilots, and deferred when it collided with the immediate crush of business and demands of client work.

The rest of the technology industry largely stayed away. Legal was difficult. The data was confidential. The language was technical. The consequences of error were serious. Workflows varied from firm to firm and often from lawyer to lawyer. Customers were famously resistant to changing how they worked. Even when a technology succeeded, the billable hour could make its efficiency economically awkward.

Legal technology was mildly interesting to the legal profession. To everyone else, it was a niche.

That era is over.

The new normal

Today, the legal industry is fast becoming one of the principal ports of the intelligence economy, a place where capital, professional expertise, software, data, and global commerce increasingly converge. The signs are everywhere.

Harvey raised $200 million in March 2026 at an $11 billion valuation.[2] Legora raised $550 million that same month at a valuation of approximately $5.55 billion, then extended the round to $600 million a month later.[3] These valuations reflect expectations about the scale of legal transformation now underway.

The deeper signal is who wants in on the action – everyone wants to get a piece of legal.[4]

Anthropic has released Claude for Legal – a suite of legal plugins, connectors, and skills designed for legal research, contract review, diligence, regulatory work, and other legal workflows.[5] Microsoft has introduced a Legal Agent directly inside Word, where lawyers can summarize, review, redline, and negotiate contracts.[6] OpenAI has long supplied foundation-model technology beneath Harvey and is extending Codex for lawyers.[7] Google, Amazon, Perplexity, and X have also entered the legal market directly.[8]

These companies are entering legal because legal work is one of the largest concentrations of expensive, language-based, knowledge-intensive activity in the economy. It is full of documents, decisions, negotiations, classifications, research, drafting, review, judgment, and process. Much of it is performed by highly compensated, licensed professionals whose work process has changed remarkably little in decades.

However, this legal technological renaissance is not only being driven from the top-down, by the largest technology companies and best-funded legal AI start-ups. It is also being supercharged from the ground up by a growing open-source community led by LegalQuants and open source platforms, such as MikeOSS, Lavern, and LQ.AI – positioned as an alternative to closed, commercial systems such as Harvey and Legora or Claude and ChatGPT.[9]

What once required a venture-backed company, a large engineering team, and years of development can increasingly be approximated by a legally trained builder working with AI coding systems and open-source components.

Open-source AI models intensify the pressure.[10] Proprietary frontier systems will continue to lead in some areas, including difficult reasoning, multimodal work, reliability, and long-horizon execution. At the same time, useful intelligence is becoming less scarce.

Now law firms and legal departments can decide which parts of the system they should own, which they should rent, and which must remain portable. The durable asset may be the software harness that surrounds the AI model, the knowledge, the workflow, the client relationship, or the architecture connecting them all together.

The system is your superpower

The first wave of LegalTech digitized the office. Word processing replaced typewriters. Email took over from letters and faxes. Filing cabinets led to document management. Online research replaced bookshelves. Practice management systems organized matters, time, billing, and clients. The second wave deployed these technologies in the cloud.

Technology made lawyers more efficient while preserving the lawyer as the central unit of production. Technology supported the work. The lawyer performed it.

Artificial intelligence changes that relationship. AI reads documents, compares positions, identifies issues, drafts language, generates alternatives, summarizes evidence, traces relationships, monitors change, and increasingly acts across connected systems. The lawyer directs, supervises, refines, and takes responsibility for the result, while the system participates directly in producing the work.

AI's participation in the work changes what organizations compete on. For law firms, the old unit of competition was largely human capacity. The firm with the strongest partners, the deepest associate bench, the best reputation, and the greatest number of available hours could command the most sophisticated work. Scale meant adding more lawyers. Leverage meant a larger pyramid. Knowledge was carried in people, transmitted through apprenticeship, and monetized one engagement at a time.

The new unit of competition is the system itself.

A firm may employ brilliant lawyers and still be structurally weak. It may buy the best AI platform and still fail to create value. It may have an innovation team, a knowledge department, a technology budget, and a responsible AI policy, yet remain less than the sum of its parts.

The market will increasingly distinguish between organizations that

possess AI tools and organizations that have learned how to orchestrate intelligence into the way they operate to create real value for their clients.

That distinction is the subject of this book.

What this means for in-house legal teams

For general counsel and corporate legal departments, the immediate promise of AI is obvious. More work can be handled internally. Contracts can be reviewed faster. Legal research can be conducted across larger bodies of information. Business teams can receive guidance without waiting in an email queue. Compliance can shift from periodic review toward continuous monitoring. Outside counsel spending, if necessary, can be examined, routed, and controlled with greater precision.

The more consequential change is structural. The legal department is moving from a service function to a legal operating system for the business.

Historically, most legal departments have worked reactively. A business unit identifies a problem, sends a request, and waits for a lawyer. AI makes another model possible. Legal judgment can be embedded into the systems where business decisions are made. Procurement software can identify unacceptable terms before an agreement reaches legal. Employment systems can surface risk before a termination becomes a claim. Corporate knowledge can become accessible at the moment of need rather than only after a lawyer is available.

General counsel must decide what work should remain human, what can be assisted, what can be automated, and what should be prevented from arising at all. They must determine which information can be trusted, which systems may access it, which models are appropriate for which data, and who is accountable when an AI-supported process influences a consequential decision. They must manage a mixed portfolio of internal capability, outside counsel, technology vendors, open-source components, alternative providers, and automated services.

They must also become more sophisticated buyers. A law firm can no longer prove its value by showing that it licensed the same AI platform as its competitors. Corporate clients will increasingly ask harder questions. What has the firm built around the model? What institutional knowledge has it structured? What outcomes have improved? What work becomes faster or unnecessary? What audit trail exists? Can the system move to another model? Who owns the workflow and the data it generates? How does the firm price the efficiency it claims to have created? Can it demonstrate that the benefit reaches the client?

The in-house leaders who ask these questions will reshape the external legal market to the benefit of their businesses.

What this means for law firms

For law firms, AI will expose weaknesses the traditional model allowed firms to ignore. A firm whose knowledge lives primarily in the heads of a few partners will discover that expertise cannot be scaled simply by purchasing a model. A firm whose systems do not connect will discover that agentic workflows cannot operate across information they cannot reach. A firm whose compensation system rewards hours will discover that training lawyers to work faster does not give them a reason to do so. A firm whose pricing remains tied to effort will discover that every efficiency gain creates pressure on revenue.

A firm that depends entirely on a vendor's AI application will also discover that application features commoditize quickly. An open-source project may reproduce the visible interface. A foundation-model company may release the capability as a plugin. The firm's durable advantage must be harder to copy. It will live in the firm's knowledge, workflows, client context, evaluation systems, trusted implementation, harness, and ability to orchestrate people, models, tools, and judgment into outcomes the client cannot easily obtain elsewhere.

The opportunity for firms that act deliberately is extraordinary. AI allows a firm's best judgment to travel further. It allows expertise developed over decades to be captured, refined, and applied across matters and outlive those who contributed it. It allows firms to design services that operate continuously rather than waiting for a client request. It makes new pricing models possible because the cost of delivery can fall while the value of the outcome remains high.

AI also lowers the cost of building. A lawyer who understands a workflow can increasingly prototype the software that executes it. A practice group can turn a playbook into an operational system. A firm can create client-facing products, automated services, managed offerings, and embedded legal capabilities that would once have required the budget and engineering organization of a software company.

Lower barriers increase the number of things an organization can build, including the number of things it should not.

The challenge for a firm is to decide what should be built, establish the conditions under which people will use it, connect it to the firm's knowledge

and systems, govern it responsibly, evaluate it continuously, and convert its use into an advantage clients can recognize.

That is a strategic question.

What this means for clients

The most important stakeholder in this transformation is the client.

For corporate clients, the opportunity is a legal service that is faster, more predictable, more closely integrated with business operations, and priced according to value rather than the amount of professional time consumed. The legal function can become less of a roadblock and source of frustration and more of an empowering enabler of business. Problems can be identified earlier. Routine work can be handled more quickly, consistently and affordably. Outside counsel can be judged by outcomes rather than activity.

For everyday people, the stakes are even greater. The traditional legal system has rationed expertise through price, geography, professional capacity, and procedural complexity. Millions of people face legal problems without meaningful assistance because the machinery for delivering legal help is too expensive, fragmented, and difficult to navigate.

AI can change the economics of justice.

Who this book is for

This book is for the people responsible for navigating this change. It is for the managing partner who knows that another pilot will not produce a strategy. It is for the general counsel who sees that the legal department's future cannot be built by adding one more application to an already fragmented stack. It is for the chief innovation officer, knowledge leader, legal operations professional, technologist, and practice leader who has been asked to make AI work inside an organization whose incentives, systems, and habits were designed for another era.

It is for the lawyer trying to understand how professional value changes when familiar tasks can be performed by machines. It is for the junior lawyer who needs a better development path than years of low-agency repetition. It is for the builder who understands that a compelling demonstration does not amount to a durable capability.

It is also for leaders in legal aid, courts, government, and public-interest organizations. The same forces changing corporate legal work can be used to increase access, extend limited expertise, and redesign services around the lives of the people who need them most.

The strategic task ahead

The market is moving too quickly for a book built around a list of products. The companies leading today will change. Valuations will rise and fall. Models will improve. Open models will narrow gaps with frontier models. Features that are new will become standard. New entrants will arrive. Some celebrated platforms will be acquired, commoditized, or forgotten. The tools available when you finish this book may be different from the tools available when I wrote it.

The underlying strategic challenge will remain.

Organizations must decide why they are changing before they decide what to buy or build. They must build the infrastructure, knowledge, trust, and human conditions that allow new capabilities to take root. They must design for model agnosticism rather than model dependence. They must treat context, evaluation, harnesses, and orchestration as architecture. They must redesign services rather than merely accelerate old tasks. They must create systems in which experience compounds instead of disappearing at the end of each engagement. They must align their economics with the value those systems produce. And they must learn continuously, because no static transformation plan can survive a technology that keeps moving.

This is happening now. Capital is moving. Clients are changing their expectations. Technology companies are entering the market. Open-source builders are challenging the assumptions beneath multibillion-dollar valuations. Legal departments are building capabilities that once belonged exclusively to outside counsel. Law firms are turning expertise into software. Individual lawyers are creating tools that would have required entire development teams a few years ago.

The boundaries between legal services, software, data, professional judgment, and organizational architecture are blurring because they are becoming parts of an integrated system.

The organizations that succeed will understand what is changing beneath the visible technology and build deliberately for that reality. They will know what they are trying to accomplish, what must be built before it can be accomplished, and why raw intelligence is only one component of capability. They will know how to transform while preserving the judgment, trust, and humanity on which legal service depends.

This book is about what, when, and how to build for our new reality. And it begins, as every sound strategy begins, with a single question.

Why.

References

1 Thomas G. Martin, "Three Countries, Three Weeks, One Truth: AI Has Arrived," LawDroid Manifesto, 19 December 2025, www.lawdroidmanifesto.com/p/threecountries-three-weeks-one-truth (quoting Richard Tromans of Artificial Lawyer).

2 Harvey, "Harvey Raises Growth Round at $11 Billion Valuation Co-led by GIC and Sequoia," 25 March 2026, www.harvey.ai/blog/harvey-raises-growth-round-at-dollar11-billion-valuation-co-led-by-gic-and-sequoia. See also Reuters, "Legal Software Firm Harvey Valued at $11 Billion in Latest Funding Round," 25 March 2026.

3 Legora, "Legora Raises $550 Million Series D to Fuel US Growth," 10 March 2026, https://legora.com/newsroom/legora-raises-550-million-series-d-to-fuel-us-growth; Legora, "Legora Extends Series D with Additional $50 Million, Welcomes Atlassian and NVentures as Investors," 30 April 2026.

4 Evan Ochsner, "Big Tech Advances Into Crowded Legal AI Arena as Next Frontier," *Bloomberg Law*, 16 July 2026, https://news.bloomberglaw.com/legal-ops-and-tech/big-tech-advances-into-crowded-legal-ai-arena-as-next-frontier. Martin is quoted in the article: "Everyone wants to get a piece of legal."

5 Anthropic, "Claude for Legal Teams," webinar, 21 April 2026, www.anthropic.com/webinars/claude-for-legal-teams; Anthropic, *Claude for the Legal Industry: A Practical Deployment Guide* (2026), www-cdn.anthropic.com/files/4zrzovbb/website/4b29cc317c727542642b5056e412cf8e779e13d8.pdf.

6 Sumit Chauhan, "Word: Legal Agent in Frontier," *Microsoft 365 Copilot Blog*, 30 April 2026, https://techcommunity.microsoft.com/blog/microsoft365copilotblog/word-legal-agent-in-frontier/4516218; Microsoft, "Get Started with the Legal Agent (Frontier)," Microsoft Support, 2026; Microsoft, "Legal Agent: Transparency Documentation," 2026.

7 OpenAI, "Harvey," customer story, https://openai.com/index/harvey/; Ryan Lopopolo, "Harness Engineering: Leveraging Codex in an Agent-First World," OpenAI, 11 February 2026.

8 Richard Tromans, "Meet Amazon Quick 'For Legal' – Updated," *Artificial Lawyer*, 15 July 2026, www.artificiallawyer.com/2026/07/15/meet-amazon-quick-for-legal; Joe Patrice, "Perplexity Jumps Into Legal With 'Computer for Counsel,'" *Above the Law*, 24 June 2026, https://abovethelaw.com/2026/06/perplexity-jumps-into-legal-with-computer-for-counsel; xAI, "AI for Legal Teams," https://x.ai/solutions/legal; Google, "AI for legal," https://workspace.google.com/intl/en_ca/resources/ai-for-legal/.

9 Richard Tromans, "Mike, the Open Source Legal AI Platform: Will Chen Interview," *Artificial Lawyer*, 4 May 2026, www.artificiallawyer.com/2026/05/04/mike-the-open-source-legal-ai-platform-will-chen-interview/; Caroline Hill, "Mike OSS: Open Source Legal AI Tool 'Changes the Negotiation,'" *Legal IT Insider*, 5 May 2026; Ken Crutchfield, "When Open Source Meets Legal: How MikeOSS Signals the End of Legal's Secret Sauce," *LawNext*, 8 May 2026.

10 Francisco Velasquez, "China's Moonshot AI Claims Kimi K3 Can Rival OpenAI and

Anthropic," BBC News, 17 July 2026, www.bbc.com/news/articles/cy9w4q8pgp0o; Esther Shittu, "Alibaba Qwen 3.8 Max Shows China Closing In on U.S. Models," *AI Business*, 21 July 2026, https://aibusiness.com/generative-ai/alibaba-qwen-3-8-max-shows-china-closing-on-u-s-models.