

Who This Book Is For

This book is written for CFOs (Chief Financial Officers), VPs of Finance, Controllers, Heads of FP&A (Financial Planning and Analysis), Treasury Directors, and anyone who leads a team and needs to make decisions on generative AI over the next twelve months.

You do not need a technical background. If you understand financial reporting, budgeting, forecasting, controls, and the realities of running a finance function, you already have the necessary context.

You also do not need to work for a bank or a large regulated organisation. The principles in this book apply to companies of all sizes. A forty-person company will not need the same governance structure as a bank with four thousand employees. However, both organisations need to select appropriate use cases, understand the risks, assign responsibility, and determine whether the technology delivers measurable value.

The examples throughout the book come from finance, but many of the decisions extend beyond it. Leaders in other functions face similar questions when prioritising use cases, evaluating vendors, establishing controls, running pilots, and managing organisational change. Finance remains the focus because it offers one of the clearest tests of whether generative AI can operate within a function that depends on accuracy, evidence, and accountability.

This book is not for data scientists or machine-learning engineers. It does not explain neural network architectures, model training procedures, or the mathematics underpinning large language models. Those subjects are covered elsewhere.

This book covers what those other books miss: how to deploy generative AI in a regulated context (e.g., finance), including all the organisational, regulatory, and human complexities involved.

What This Book Will Help You Do

This book will help you distinguish useful applications of generative AI from ideas that are impressive in a demonstration but unreliable in practice.

It begins by **explaining how generative AI is changing finance work**. This change is not simply another wave of automation. Traditional automation is effective when a process follows stable rules and must produce the same result every time. Generative AI is useful when the work involves language, interpretation, investigation, or communication. Understanding that distinction is essential, as many disappointing AI projects begin with the wrong tool for the task.

You will learn how **deterministic systems, conventional automation, generative AI, and AI agents** fit together. The objective is not to replace one approach with another, but to use each where it performs best. Calculations, reconciliations, validation rules, and financial controls often require deterministic methods. Drafting, summarisation, document analysis, exception investigation, and commentary may benefit from generative AI. Agentic approaches can coordinate multiple steps, but they also introduce additional cost, risk, and assurance requirements.

The book then maps the **most credible applications** across finance. These include management commentary, financial communication, meeting summaries, policy questions, data quality narratives, accounts payable exceptions, contract review, forecast analysis, intercompany reconciliation, cash collection, provision documentation, anomaly investigation, Internal Audit support, and regulatory change analysis.

Not every use case warrants the same level of attention or investment. The **Use Case Maturity Matrix** distinguishes Quick Wins, Strategic Bets, Learning Pilots and Defer/Stop applications. The prioritisation framework then helps you compare opportunities by value, feasibility, data readiness and risk.

Once you have identified a promising use case, the book helps you **assess whether your organisation is ready to pursue it**. The readiness assessment examines the conditions surrounding the project, while the Post-Assessment Action Matrix outlines what to address before implementation. The Three Horizon Approach places the first pilot within a broader path from initial evidence to integration and eventual scale.

You will also learn how to select **the right first team**, address common executive objections, and build a credible business case. The first project aims not to transform the entire finance function. It is to produce useful evidence, build organisational confidence, and earn the opportunity to pursue a second project.

The book **examines the mistakes** that most often prevent this. These include trusting polished output, starting with an unnecessarily difficult problem, buying a tool before understanding the process, neglecting change management, postponing governance, measuring the wrong outcomes, expecting one model to suit every task, and abandoning the initiative when the first results are imperfect.

You will learn how to decide whether to **use, build, buy, or develop a solution with a partner**. The decision depends on the use case, data sensitivity, your team's capabilities, required speed, and the level of control you need. The book also examines vendor evaluation, contract terms, model choice, operating costs, and the risk of becoming dependent on a single provider.

Data, governance, and risk receive detailed treatment because they determine whether an AI initiative can move beyond experimentation. The book explains how to assess data

readiness, set access rules, define review responsibilities, maintain logs and audit trails, manage incidents, and develop a practical AI policy.

It also examines the **regulatory environment across Europe, the United States, the United Kingdom, Singapore, and the wider Asia-Pacific region**. The objective is not to turn finance leaders into lawyers. It identifies which questions to ask, which responsibilities cannot be delegated to a vendor, and how to extend existing control frameworks to cover AI.

The **pilot chapter** turns these ideas into an **eight-week implementation process**. It covers preparation, historical testing, live operations, evaluation, cost, risk, and leadership communication. A worked example follows a variance commentary use case from initial setup through to operational adoption. The chapter also provides a practical method for writing instructions and prompts that are clear enough for finance work.

Technology alone does not drive adoption. The book therefore examines the **skills that finance leaders and their teams now need**. These include AI literacy, critical review, data fluency, process design, change leadership, vendor management, and ethical judgement. It also considers emerging responsibilities, such as the AI champion, Data Steward, Governance Lead, and Process Redesign Specialist, and how those responsibilities can be combined within a smaller team.

The **change-management playbook** addresses the human questions that are often overlooked. Employees may fear job losses, distrust the tools' quality, worry about appearing incompetent, or stay attached to familiar processes. Training alone cannot resolve these concerns. People need to understand why the change is happening, see how it affects their work, build confidence through practice, and receive ongoing support after the pilot ends.

The final chapter brings everything together in a **90-day action plan**. It guides you from initial orientation through preparation, execution, evaluation, and deciding what should happen next. The Adoption Lifecycle then places those first 90 days within the longer journey from isolated experiments to consistent use across the finance function.

By the end of the book, you should be able to identify where generative AI can create value, select an appropriate delivery approach, establish proportionate controls, run a meaningful pilot, prepare your team, and define the next practical step for your organisation.