THE EU AI ACT PLAYBOOK

For Funds and Financial Services

A Practical Guide to AI Governance and Compliance

Shane Brett, PhD

THE EU AI ACT PLAYBOOK

FOR FUNDS AND FINANCIAL SERVICES

A Practical Guide to AI Governance and Compliance

The Practical AI Governance Library for Financial Services

Shane Brett, PhD

TABLE OF CONTENTS

CHAPTERS

APPENDICES

Introduction

Why I Wrote This Book

There is a moment that happens in every major regulatory cycle in financial services. I have seen it play out with AIFMD, with GDPR, with MiFID II, with DORA. A new regulation lands. The lawyers write the summaries. The conferences fill up with panels. Compliance teams read the briefings, attend the webinars, and add something to the to-do list.

And then nothing happens. Because nobody has written the practical guide.

Not the legal analysis. Not the policy paper. The actual guide — what does this mean for my firm, what do I specifically need to do, and what order do I do it in?

This is that book.

What Clients Actually Needed

When firms I was working with started asking me about the EU AI Act, I went looking for a practical guide, written for a compliance officer or COO who needed to understand what they actually had to do — in plain English, in their firm, on Monday morning.

I could not find it.

Everything I found was written by lawyers for lawyers. Dense, cautious, full of qualifications. Or it was so generic it was useless for anyone in financial services and funds trying to understand their specific obligations.

Meanwhile, the firms asking me questions were not abstract entities. They were fund administrators in Dublin running AML screening tools they had never classified. They were UCITS managers using suitability assessment tools they had never assessed for AI Act exposure. They were US asset managers with European investors who had never heard of the regulation and assumed it was someone else's problem.

It was not someone else's problem. It was theirs. And they needed help.

This book fills that gap.

Why Ireland Is at the Centre of This

This book focuses significantly on the Irish regulatory landscape, and it is worth explaining why.

Ireland is the second largest fund domicile in Europe and the third largest in the world[1], with over €5 trillion in assets under management in Irish-domiciled funds. Every major US and European asset manager of consequence either has an Irish-domiciled fund, an Irish management company, or both. The Central Bank of Ireland is, for practical purposes, one of the most important financial regulators in the world for any firm operating in the global funds industry.

Ireland is also uniquely positioned in the AI landscape more broadly. OpenAI and Anthropic — the two companies whose technology is driving the AI revolution — both have their European headquarters in Dublin. Google, Meta, LinkedIn, TikTok and Apple all have their primary European base of operations in Ireland.

In April 2026, the Washington Post reported that Ireland is positioning itself as the primary gateway for American AI companies expanding into Europe[2] — a designation that reflects not just existing infrastructure, but deliberate government strategy. Anthropic recently announced it is expanding its Dublin office sixfold and hiring 200 additional staff.

In July 2026, Ireland assumes the rotating presidency of the European Council — giving it direct influence over EU AI policy at exactly the moment the regulatory framework is being finalised and embedded.

Ireland is not a peripheral jurisdiction for AI regulation. Its dual nexus of global technology giants and fund companies positions it uniquely at the heart of EU AI governance. When the Central Bank of Ireland starts enforcing the EU AI Act — and it will — the impact will be felt in New York, Boston, Singapore, and Sydney, not just in Dublin.

---

[1] Irish Funds (2025) *Why Ireland 2025*. Available at: https://www.irishfunds.ie/news-knowledge/news/now-available-why-ireland-2025-publication/

[2] The Washington Post Intelligence (2026) *Ireland is positioning itself as the industry-friendly foothold for American AI companies in Europe*. AI & Tech Brief, April. Available at: https://www.washingtonpost.com/wp-intelligence

The Regulators Here Are Not Shy

There is one thing about Ireland that every non-European firm needs to understand before reading anything else.

The regulators here have demonstrated, over the past decade, that they are willing to **issue very large fines to very large companies. Including American ones.**

The numbers are worth knowing.

The Irish Data Protection Commission — Ireland's GDPR regulator — has issued some of the largest regulatory fines in European history, almost all of them against US technology companies:

- **Meta: €1.2 billion** (May 2023) — for unlawfully transferring European user data to the United States. The largest GDPR fine in history, anywhere in the world.
- **TikTok: €345 million** (September 2023) — for failures in handling minors' data.
- **LinkedIn (Microsoft): €310 million** (October 2024) — for misuse of data in its advertising systems.
- **Meta: €265 million** (November 2022) — for failures in data protection by design.
- **Meta: €251 million** (December 2024) — following a data breach affecting millions of European users.
- **Meta: €91 million** (September 2024) — for storing user passwords in plain text.
- **Meta (WhatsApp): €225 million** (September 2021) — for transparency failures.

These are not small companies that got caught out. These are the largest technology businesses in the world, with armies of lawyers and compliance professionals. They still got fined. By a regulator based in Dublin.

The Central Bank of Ireland operates with the same approach. In November 2025, it fined Coinbase Europe **€21.5 million** for failures in its AML transaction monitoring — gaps that left hundreds of millions of euros in transactions unreviewed for years. Coinbase is an American company. Operating through an Irish entity. It did not insulate them from anything.

The EU AI Act gives these same regulators new enforcement powers and new grounds to act. By December 2027, the Central Bank will have the authority to require any regulated financial services firm to demonstrate that its AI systems are properly governed, documented, and overseen. Firms that cannot **demonstrate this**

face penalties of up to **€15 million or three per cent of global annual turnover — rising to as much as seven per cent of global annual turnover for the most serious breaches** — whichever is higher.

If you are reading this from New York, Boston, London, or the Cayman Islands, those numbers are not someone else's problem.

Who This Book Is For

This book is written for three types of reader:

1. **The Irish or European fund and financial services firms** that knows the EU AI Act is coming and needs to understand what it actually requires. Fund managers, fund administrators, AIFMs, UCITS management companies, wealth managers, banks, and insurers operating in Ireland and across the EU. You are the primary audience. Every chapter is written with your specific operational reality in mind.

2. **The non-European fund manager or financial institution** that has European investors, European clients, or a European-domiciled structure. The EU AI Act applies to you regardless of where you are based. If your AI systems affect people in the EU — and for most global managers, they do — this regulation already applies. **Chapter 3** explains the mechanics. The rest of the book tells you what to do.

3. **The compliance professional, COO, or board member** who needs to understand this regulation at a level sufficient to make decisions and ask the right questions — without becoming a legal expert in EU AI regulation. This book is also for you. It will not tell you every nuance of every article. It will tell you what the Act means for a firm like yours, what you are required to do, and what the consequences are if you do not.

What This Book Is Not

This book is not legal advice. It is a practical guide written by a practitioner with deep experience in investment funds, financial services and AI governance. For specific legal questions, engage a qualified legal adviser. What this book will do is give you the foundation to have that conversation intelligently — and to avoid spending expensive legal time on questions you could have answered yourself.

This book also does not attempt to cover every single aspect of the EU AI Act. It focuses on the provisions most relevant to financial services and asset managers firms in their capacity as Deployers of AI systems — which is where the vast majority of firms sit.

A Note on Timing

This book was written and published in April 2026. The EU AI Act and its associated legislative developments were moving quickly at the time of writing. The key deadlines and regulatory facts in this book reflect the position as of that date and will be **updated quarterly** in the book over the next two years. The updated versions will be provided for free (in pdf format) to all book purchasers

The regulation will continue to evolve. Technical standards will be published. The Central Bank of Ireland will issue guidance. The Digital Omnibus will be formally adopted. Where possible I have flagged areas where further clarification is expected.

What will not change is the direction of travel. The core requirements of the Act — document your AI systems, classify them by risk, implement human oversight, maintain evidence of compliance — are not going to be amended away. The work this book describes is the work that needs to be done, regardless of the final legislative text.

If you are waiting for perfect clarity before starting, you will still be waiting in December 2027.

Start now.

My Background

I have spent over twenty-five years in the global investment funds industry, most of it working in compliance, governance and FinTech. I have worked through the implementation of AIFMD, UCITS V, MiFID II, GDPR, and DORA. I know what it looks like inside a compliance team when a major new regulation lands. I know the confusion, the waiting for clarity, and then the sudden scramble when the deadline gets close.

I have watched this happen with the EU AI Act already.

I also founded and scaled a leading US VC-backed governance technology company that built compliance management software for the global funds industry. We raised venture capital from institutional investors in the US and Europe, scaled the business globally, and navigated exactly the kind of compliance challenge this book addresses — from both sides. Both as a firm subject to regulation, and as a technology Provider helping firms manage it.

After completing a PhD in innovation and an Advanced Diploma in AI, I became genuinely fascinated by the governance challenge that AI creates for financial services and fund management firms. The regulatory frameworks were arriving fast, the practical guidance was nowhere to be found, and many of the firms that needed help most had no idea where to start. So, I built a team and a consultancy specifically to fix that.

I am not writing this from a theoretical perspective. I am writing it from the perspective of someone who has done this work, in this industry, in this regulatory environment, and who knows what is useful and what is not.

This is practical guidance, grounded in real-world experience of this industry. If you need help applying it, get in touch.

Shane Brett, PhD, Global Perspectives Dublin, April 2026

CHAPTER 3 - Does This Apply to You? Mapping Your Role

The Question Everyone Gets Wrong First

When financial services firms first encounter the EU AI Act, the first question is usually: "Does this apply to us?"

That is the wrong question.

The right question is: "**In what capacity does it apply to us** — and to which of our systems?"

Your obligations under the Act depend entirely on what role you play in relation to any given AI system. The same firm can simultaneously be a "**Provider"** for some tools, a **"Deployer"** for others, and **outside the scope** for others entirely. Getting this wrong costs money. Treat yourself as a pure Deployer when you are actually a Provider, and you will underinvest in compliance. Treat yourself as a Provider when you are a Deployer, and you will overinvest. Both outcomes are expensive.

The Three Roles That Matter

1. **Provider**

You built it. You deploy it under your own name. You are responsible for it from the ground up — technical documentation, conformity assessments, risk management systems, registration.

Most financial services firms are **not** Providers in the classic sense. But if your firm has developed **its own proprietary risk model**, its own algorithmic trading system, or its own AI-powered client analytics tool — and you deploy that under **your own name** — you are a **Provider** for those systems. The obligations are heavier.

The Provider definition has not changed materially under the Digital Omnibus. The core test remains: did your firm develop the system and place it on the market or put it into service under your own name? If yes, you are a Provider. If you are in any doubt about your classification, get specific legal advice before assuming you are a Deployer.

2. Deployer

You use an AI system built by someone else in the course of your professional activity. This is where the **vast majority** of financial services firms sit.

Deployer obligations are less extensive than Provider obligations — but they are not light. **Article 26** of the Act sets out what Deployers must do. **Chapter 5** covers those obligations in detail.

3. Importer

You are an EU-established firm that places on the EU market an AI system built by a Provider located outside the EU. This role is more common in financial services than most firms realise.

If your firm is the Irish or Luxembourg entity through which a US-built AI platform is made available to EU clients or investors — whether that is an AML screening tool, a KYC platform, or a portfolio analytics system built by a non-EU vendor — you may be acting as an importer under the Act.

Under **Article 23** Importers carry **specific obligations,** which include:

- verifying that the non-EU Provider has met their obligations,
- ensuring the Provider has technical documentation in place,
- and confirming the system bears the required conformity markings before it is deployed.

In practice, this means your vendor due diligence process takes on additional weight. If your US AI vendor has not met their EU AI Act obligations as Provider, your firm — as the EU importer — **is exposed**.

You cannot simply pass that responsibility back across the Atlantic.

The Provider-Deployer Trap

Here is where firms get into trouble.

The boundary between Provider and Deployer is not always clean. If your firm takes a third-party AI platform and **substantially customises it** — adding your own data feeds, retraining elements of the model on your own data, adjusting outputs to produce your own branded risk assessments — are you a Deployer? Or, by virtue of that modification, have you become a Provider?

The more substantially you modify a third-party system, the more it starts to look like a Provider relationship. If you are in this position, get specific legal advice. Do not assume. The wrong assumption in this direction is the expensive one.

This boundary was flagged for clarification during the Digital Omnibus negotiations in early 2026, but no definitive legislative fix has yet been adopted. The practical test remains: the more you change how the system works — not just how it is configured — the more you look like a Provider. When in doubt, treat yourself as one and take professional advice.

Vendor Responsibility — The Part Firms Consistently Get Wrong

As a Deployer, you are responsible for ensuring that the AI systems you use comply with the Act. You cannot outsource that responsibility to the vendor. This applies equally to EU-based firms and non-EU firms. If your AI vendor is non-compliant, your firm is still exposed — regardless of where either of you is based.

This surprises people. They assume that if they are buying a product from a reputable vendor, compliance is the vendor's problem. It is not. If the vendor's system is non-compliant, your firm is still exposed.

Think of it this way. Under financial services regulation, you cannot delegate regulatory responsibility to a service Provider and then walk away. DORA makes this explicit for ICT risk. The EU AI Act makes it explicit for AI systems. You are responsible for knowing what your AI tools do, how they are governed, and whether they meet the Act's requirements — even if someone else built them.

This has two immediate practical implications.

Firstly, start asking your AI vendors the right questions.

1. What is the risk classification of your system under the EU AI Act?
2. What documentation do you maintain?
3. How do you support Deployers in meeting their Article 26 obligations?

If your vendor cannot answer these questions clearly and completely, that is itself important information about the risk in that relationship.

Secondly, review your vendor contracts.

The Act creates specific obligations around what Providers must make available to Deployers. If your current vendor agreements do not address any of this, they need to be updated. **Chapter 5 covers** what those contracts should contain. **Appendix E** contains a vendor due diligence questionnaire you can use with your AI vendors today.

US & Non-EU Funds and Financial Services Companies.

If you are based in the US, the UK, Cayman, Asia or anywhere outside the EU, pay close attention here.

The EU AI Act applies to you.

Not as a technicality. Not as a theoretical future risk. Right now, as a real and operational obligation. The Act applies to any Provider or Deployer whose AI systems **affect people located in the European Union** — regardless of where you are established.

And yes — it applies in full. Not in a watered-down, reduced version for non-EU firms. If your AI systems affect EU persons, the same obligations that apply to a Dublin or Paris-based firm apply to you. The only practical difference is the enforcement mechanism, which — as Chapter 1 explained — will reach you through your LP base before it reaches you through a regulator.

What "affects EU persons" means in practice

If your fund has European investors — Irish, Luxembourg, Dutch, German — and your AI systems are involved in managing their capital, screening their onboarding, assessing their suitability, or monitoring transactions involving their assets — the EU AI Act applies to those systems.

Let me make that concrete with examples from across financial services.

- **A New York hedge fund** with Irish and Luxembourg investors uses an AI-powered investor onboarding tool for KYC verification. When those European investors go through onboarding, that AI system is affecting people located in the EU. The Act applies. Full stop.
- **A Boston asset manager** runs a quantitative strategy using an AI-driven credit scoring model. The fund has European pension fund clients. The outputs of that credit scoring system are affecting the financial interests of European investors. High-risk provisions under the Act are engaged.
- **A Cayman-domiciled fund** markets to European institutional investors under AIFMD's national private placement regime. The investment process includes an AI system making automated recommendations on position sizing. European investors' capital is being managed, in part, by an AI system. The AI Act applies.
- **A US bank** with European retail or corporate banking clients uses an AI-driven fraud detection and transaction monitoring system. Those systems are generating risk scores and flagging transactions involving EU persons. High-risk provisions apply.
- **A US insurance company** uses AI-driven underwriting and risk pricing models for European policyholders. Annex III explicitly lists life and health insurance AI pricing systems as high risk. If your model’s price or underwrite European risks, you are in scope.

Why This Will Hit US Firms Before Regulators Do

Here is the part that most US managers are not thinking about yet.

The EU AI Act will not reach across the Atlantic and fine a New York firm in 2028. That is not the near-term enforcement mechanism.

The enforcement mechanism is **your investor base.**

European institutional investors — pension funds, insurance companies, sovereign wealth funds, family offices — are already beginning to add AI governance questions to their standard due diligence frameworks. Not because a regulator told them to. Because their own boards, their own risk committees, and their own legal teams are asking about AI governance in the assets they own.

This will cascade fast. Over the next year, AI governance questions will become standard in LP DDQs across the European institutional market. Firms that cannot answer them will not be allocated capital. Not because they are penalised. Because they will be perceived as being behind the curve — and in a competitive fundraising environment, that is enough.

European managers who build their AI governance frameworks now will be ahead of their US counterparts. Not slightly ahead — meaningfully ahead, in a way that shows up in due diligence scores, allocation decisions, and investor conversations.

US managers who have not started will still be building their frameworks while European competitors are already demonstrating them. That window of competitive advantage is available right now. It will not stay open.

European managers who build their AI governance frameworks now will be able to demonstrate that readiness to global investors. US managers who have not started will be answering DDQ questions with "we are working on it”. This is not an answer that institutional investors will accept.

EU & Non- EU firms - A Practical Framework for Determining Your Role

These five questions apply to every firm — EU-based or not, fund manager or bank, large or small. For any AI system your firm uses, ask them in order.

Question 1: Did we build it?

If yes, and you deploy it under your own name, you are almost certainly a Provider.

Question 2: Did we substantially modify it?

If you took a third-party system and meaningfully changed how it works — not just configured it — you may have become a Provider. Get legal advice.

Question 3: Are we using it professionally?

If yes, and you did not build it, you are a Deployer. Review your Article 26 obligations in **Chapter 5**.

Question 4: Does it affect people in the EU?

If your business involves European clients, investors, or employees, the Act applies regardless of where you are located.

Whether AI systems used in marketing or prospecting activities directed at potential European investors also trigger the Act is an emerging question –one to watch as guidance develops, and one worth raising with your legal and governance advisers now.

Question 5: Is it High Risk?

If the system makes or influences decisions about creditworthiness, financial access, investment management, or AML screening, the answer is almost certainly yes. And if the answer is yes, there is a substantial body of work to do: documentation, risk classification, human oversight frameworks, performance monitoring, vendor reviews, log maintenance, and client disclosures. **Chapter 5** covers all of it.

Article 22: The Authorised Representative Requirement

There is one more important potential obligation Non-EU Providers cannot ignore.

If your firm is based outside the EU and has developed its own high-risk AI system — a proprietary risk model, a quant signal engine, a KYC scoring tool built in-house — there is an additional requirement beyond everything I have described above.

Under Article 22 of the EU AI Act, you must **appoint an EU-based authorised representative** by written mandate before making that system available to users in the European Union. This is not optional, and it is not something that can be addressed retrospectively.

Your authorised representative becomes the point of contact for EU regulators, holds your compliance documentation, and is empowered to cooperate with competent authorities on your behalf. Critically, they must retain key documentation — including technical records and the EU declaration of conformity — for ten years after the system is placed on the EU market.

This is not a paperwork formality. If an EU regulator wants information about your AI system, they go to your authorised representative. And if that representative concludes you are not meeting your obligations under the Act, they are required to terminate the mandate and notify the relevant authorities immediately.

Note also that the authorised representative role carries real legal responsibility. It should be held by a qualified EU-based legal entity — a law firm or specialist service provider — not simply an existing office or sales function.

The practical step: If you have built proprietary AI systems used in connection with European clients or investors, confirm whether Article 22 applies and identify a qualified EU-based entity to serve as your authorised representative. Add this to your Phase 1 AI inventory work.

Chapter Summary

- Your obligations depend on your role: Provider, Deployer or Importer.
- Most financial services firms are Deployers. Firms that substantially modify AI systems may be Providers. When in doubt, take legal advice.
- The Provider-Deployer boundary has not been definitively clarified by the Digital Omnibus. The practical test remains: the more you change how a system works, the more you look like a Provider.
- If you are a non-EU firm with European clients or investors, the EU AI Act applies to you. In full. Not in a reduced version. The industry awaits clarification regarding the status of marketing to potential European investors.
- For Non-EU firms, this will show up in LP and investor due diligence before regulators act. That is the real commercial deadline.
- European managers who act now have a competitive window that will not stay open.
- As a Deployer, you cannot outsource AI Act compliance to your vendors. You are responsible for understanding and evidencing compliance across your entire AI stack.

CHAPTER 4 - High-Risk AI in Financial Services: Are You Already Running It?

You Are. Almost Certainly.

Let me ask you something.

- Does your firm use software to screen clients against AML watchlists?
- Does it use any automated tool to assess investor suitability or creditworthiness?
- Does it run any AI-driven process that contributes to a decision about who gets access to a financial product or service?

If you answered **yes** to any of those questions, **you are already running high-risk AI** under the EU AI Act.

Not "might be." Not "possibly." You are.

And the governance framework that the Act requires around those tools almost certainly does not exist in your firm yet.

You are not alone. The vast majority of Irish and European financial services firms — and US, Asian and global firms with EU exposure — including fund managers, fund administrators, insurers, banks, AIFMs and UCITS management companies — are in exactly the same position. They have been using AI-driven tools for years. They have never classified those tools under any regulatory framework. They have no documentation, no oversight logs, no risk management system around them.

From December 2027, the Central Bank of Ireland and other European regulators will have the power to ask for all of it. And they will use it. Non-compliance with the high-risk AI provisions carries penalties of up to €15/€35 million or three/seven per cent of global annual turnover — and the regulators enforcing these rules have a documented track record of acting.

European national financial regulators imposed over €100 million in financial services fines in 2024 alone. Germany's BaFin fined Citigroup €12.9 million in 2024 for algorithmic trading control failures under MiFID II. These are not cautionary tales from another era. They are the last few years' headlines.

A note on the UK. The UK has not adopted the EU AI Act and is taking a principles-based, sector-led approach with no binding AI law yet in force. That does not mean UK firms are off the hook. Any UK financial services firm with EU clients, investors, or operations (i.e. most of them) is subject to the EU AI Act regardless of where it is headquartered.