AI Regulation and Litigation

Carolyn Saund, PhD

carolynsaund@gmail.com
http://www.carolynsaund.me

Quick Intro

- Computer Science undergrad, SWE at several startups in robotics and affective AI in the US for a few years
- PhD in Computer Science Glasgow 2022, working in a technical and economic litigation consulting firm since then
- Background and lens is sorta always Cognitive Science – why and how do people make the decisions the way they do?
- 31 Weeks

Agenda

- AI Regulation
  - Recognize Drivers and Implications of Regulation
  - Global Regulatory Approaches to AI
  - Deep dive into UK and EU frameworks

Pee Break

- AI Litigation
  - How litigation and regulation are complementary
  - Map litigation and technical challenges
  - Case studies and discussion

AI Regulation

Why Regulate AI?

Trust and Adoption

- Consumers adopt novel products when risks are bounded and remedies exist (Data and Marketing Association (DMA))

Market Integrity

- Prevent competitive harm to markets (Competition and Markets Authority (CMA))
- Facilitate investment and growth of AI as a tool

International Geosecurity

- Industrial policy, national security, and strategic advantage (AI Security Institute (AISI))

Safety, Fairness, Values

- “we live in a society, man”
- Prevent harmful externalities that we as a society agree are bad (bias, privacy harms, unsafe autonomy)

Why Regulate AI?

Trust and Adoption

- Consumers adopt novel products when risks are bounded and remedies exist (Data and Marketing Association (DMA))

Market Integrity

- Prevent competitive harm to markets (Competition and Markets Authority (CMA))
- Facilitate investment and growth of AI as a tool

International Geosecurity

- Industrial policy, national security, and strategic advantage (AI Security Institute (AISI))

Safety, Fairness, Values

- “we live in a society, man”
- Prevent harmful externalities that we as a society agree are bad (bias, privacy harms, unsafe autonomy)

Trust and Adoption: Consumer Safety Edition

“Regulation exists to facilitate innovation”

- Compliance as market access cost that yields trust
- Consumers adopt novel products when risks are bounded and remedies exist (Data and Marketing Association (DMA))
- Certification & assurance
  - Conformity assessments lower buyer risk
- Avoid race-to-the-bottom
  - Guardrails prevent negative-sum competition

Market Integrity

- Prevent competitive harm to markets (Competition and Markets Authority (CMA))
- Facilitate investment and growth of AI as a tool without monopolization

- Examples: blocked merger btw OpenAI/Microsoft, blocked investments to Anthropic, etc

Predictable rules → outside investment

1. Global private AI investment hits record high with 26% growth.

Corporate AI investment reached $252.3 billion in 2024, with private investment climbing 44.5% and mergers and acquisitions up 12.1% from the previous year. The sector has experienced dramatic expansion over the past decade, with total investment growing more than thirteenfold since 2014.

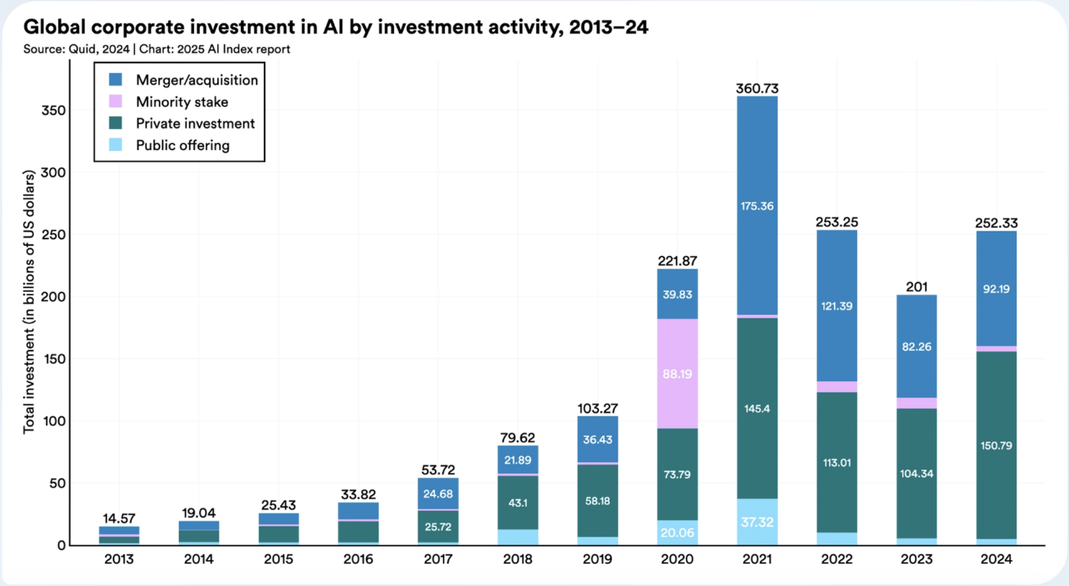

Human-Centered Artificial Intelligence Institute Index Report, Stanford University, 2025

Cybersecurity and Geopolitics

- If one of us knows, everyone knows
  - Compare to cybersecurity field at large and “black hat” vs. “white hat” hackers to “red teams”
  - Debate often breaks down to open-source vs. closed-source models, weights, and algorithms.

deepseek

-transparency isn’t always ideal, but there are good arguments for both sides – reasonable minds can differ on this

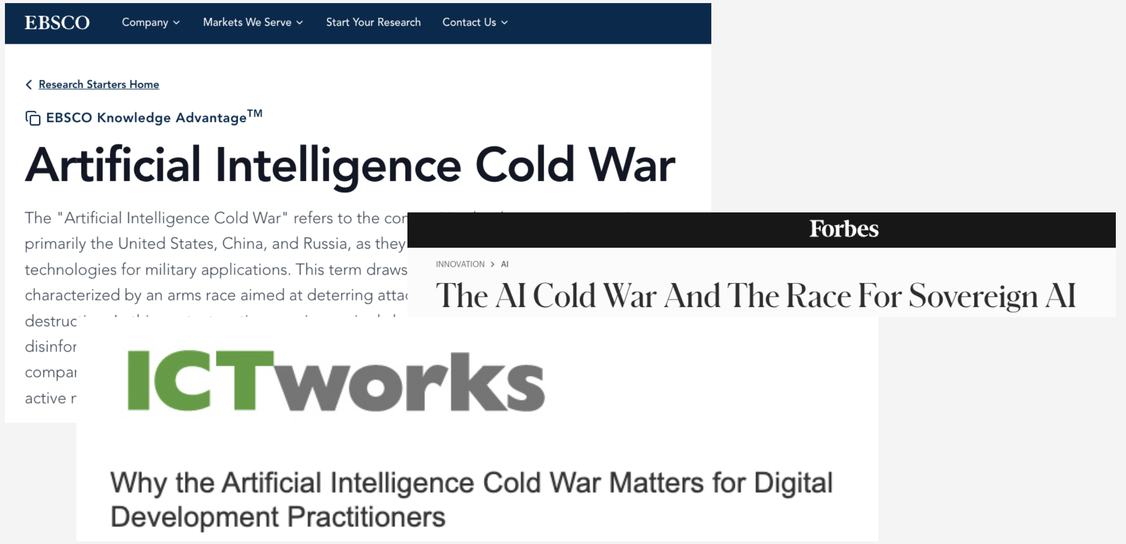

  - The AI Cold War is happening now
    - The reality is with integrated technology and evolution of an internet-based society, we are one bad cyberattack away from societal shutdown
    - This will only increase as SmartGrids come online as well.

co op

Regulation enforces social values

Values → Intentions → Enforcement

- As a society we have things we want to be true
- REGULATION IS HERE
  - we create intentions to facilitate everybody following those values
- LITIGATION IS HERE
  - Can those intentions be reasonably enforced?

- Ex. Artist copyright enforcement – “artist should be paid for their work because we as a society value art”
  - “you have to pay artists when you use their work”
  - What does it mean to “use” a piece of work in the context of training ML models?
    - It’s certainly not specifically enjoyed as a piece of art
    - It’s unlikely to be replicated
    - But it still feels wrong that

AI Regulation is especially challenging

- Very, very few people understand the implications of AI regulation
  - So new and vast, different people come at it from wildly different angles...
  - Geo-cyber-security, military defense, competition, social good, consumer products, research progress
  - AI Anxiety prevents curiosity

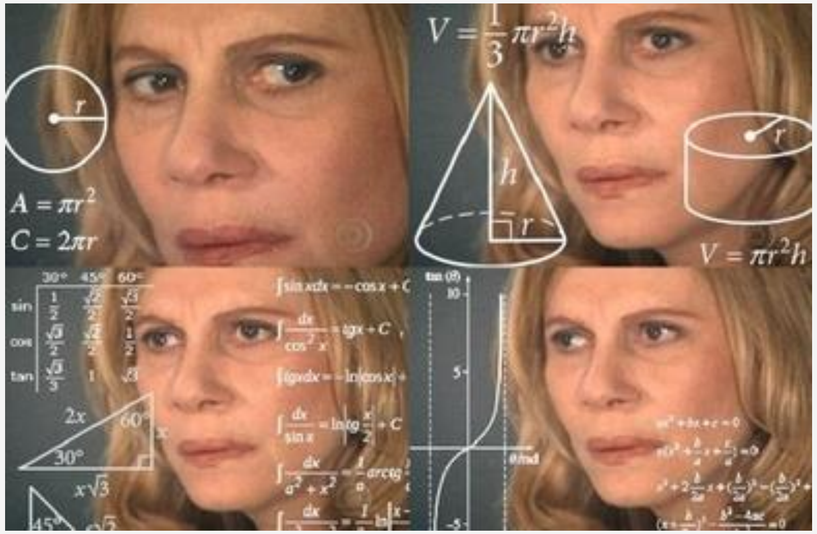

- Even fewer understand how to enforce it
  - Audits? Tech committees?
  - Where are these people coming from? The ratio of people who understand how ML models work to people who want to work in AI regulation (vs. work on the tools / research themselves where they get paid big bucks) is not on the side of government salaries let me tell you

- And companies are prepared to take advantage of that.

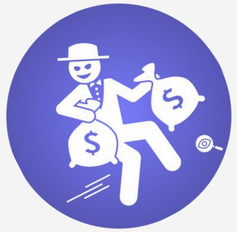

So what is regulation in practice?

ofqual

Rules about implementations	Guidelines on how to follow those rules	Audits and investigations
• Laws (EU AI Act) require bias audit, documentation, human oversight for AI hiring systems. What a company must do prior to deployment	• Frameworks like NIST, ISO to shape risk assessments and bias testing requirements. “Reasonable care”	• Regulators (EEOC, EU authority) review documentation, internal systems incl. training data, validate results. Sometimes external audits
Fines, orders, recalls	**Public pressure and socializing norms**	**Litigation**
• Financial penalties, suspending system for remediation. EU AI Committee, European Commission, AISI in UK	• Media, advocacy groups, customers react to bias reports; consumer pressure pushes reform in conjunction with regulatory risk	• Applicants file discrimination law suits, which triggers discovery into model workings and governance. Court interprets compliance efforts. Everybody is sad.

Regulation is not an event, it’s an ecosystem

AI Regulation is a Global Patchwork

North America
- United States
  - AI Regulations, Federal Trade Commission
  - AI Principles, The US DoD
  - AI RMF, NIST
- Canada
  - Artificial Intelligence and Data Act

Europe
- European Union
  - EU AI Act
- United Kingdom
  - A pro-innovation approach to AI Regulation

Asia
- China
  - Interim Measures for the Management of Generative Artificial Intelligence Services
- Japan
  - AI Guidelines for Business
- Singapore
  - The Model AI Governance Framework

Middle East
- Saudi Arabia
  - AI Ethics Principles

South America
- Brazil
  - Proposed AI Regulations

Africa
- NA

Australia & Oceania
- Australia
  - AI Ethics Principles

- Common threads
  - Transparency, risk mgmt., accountability – key things that frameworks basically set out to provide.
- Key differences
  - Remember those values we talked about?
  - Scope, upstream/downstream, penalties
- Strategic implication:
  - Multijurisdictional compliance during design is essential

EU AI Act Risk Tiers

Definitions

Risk-Tiered Enforcement Model

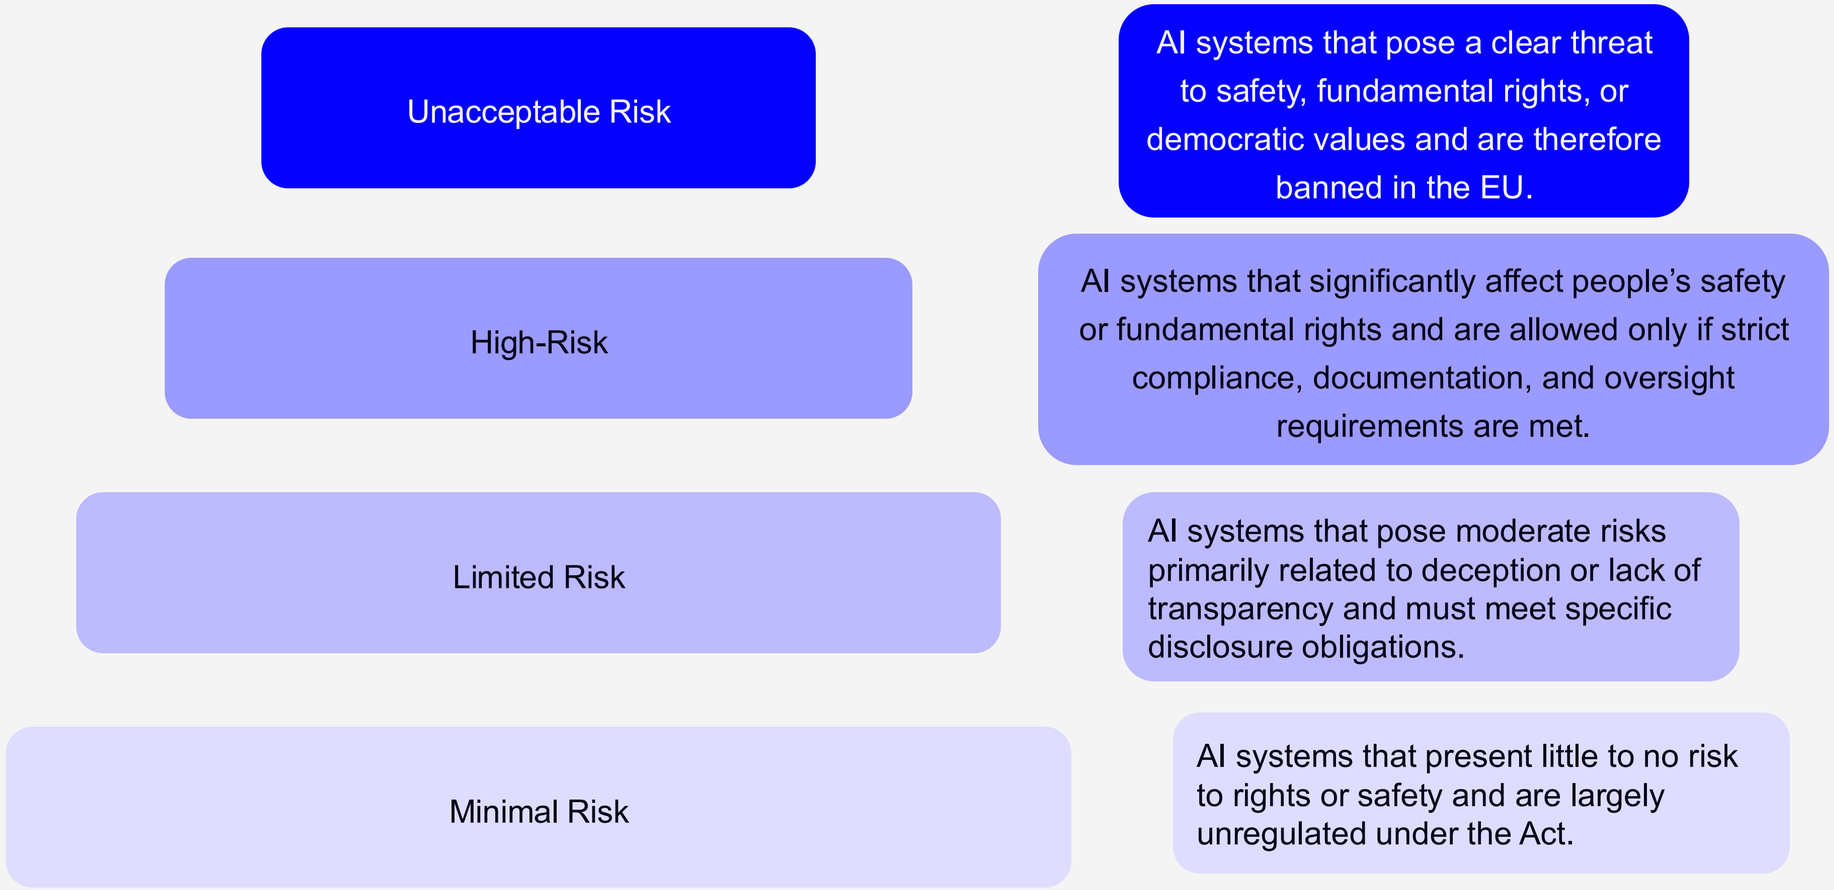

EU AI Act Applications

But also you have different model types

Risk-Tiered Enforcement Model

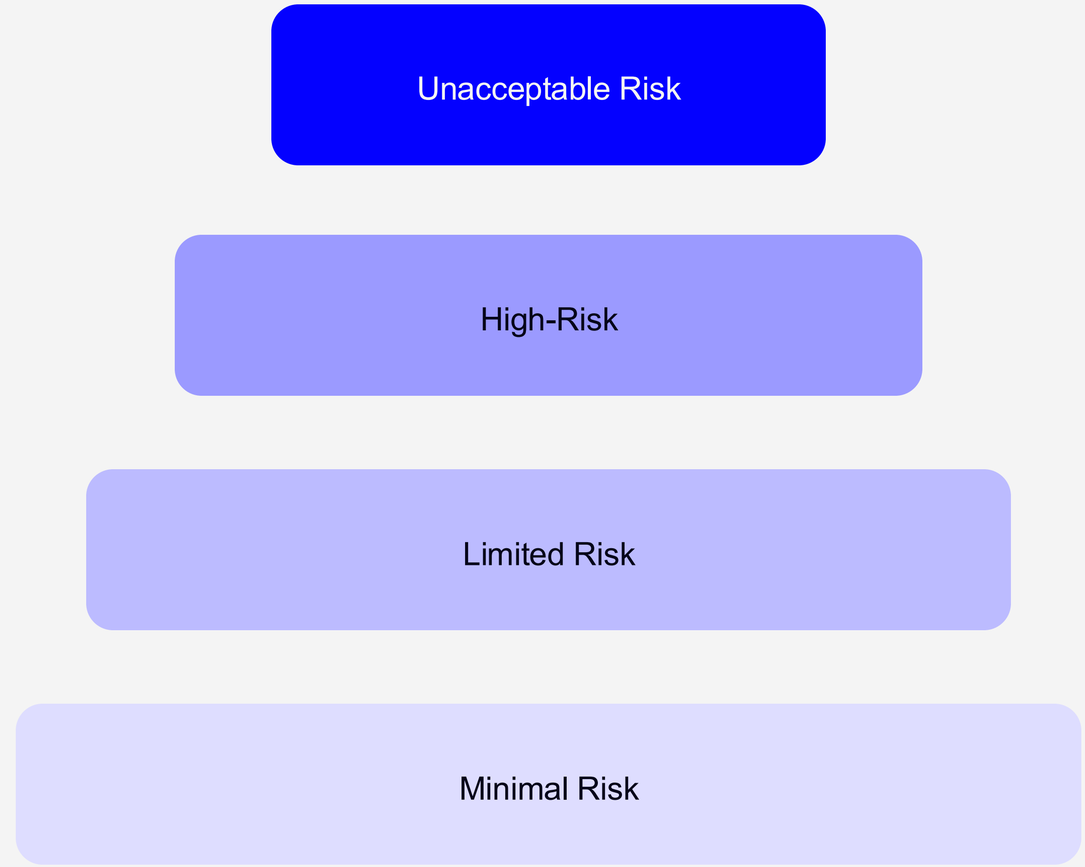

Downstream Integrators

Application-level obligations

GPAI / Foundation Models

Baseline transparency in training, documentation

EU AI Act Applications

But also you have different model types

Risk-Tiered Enforcement Model

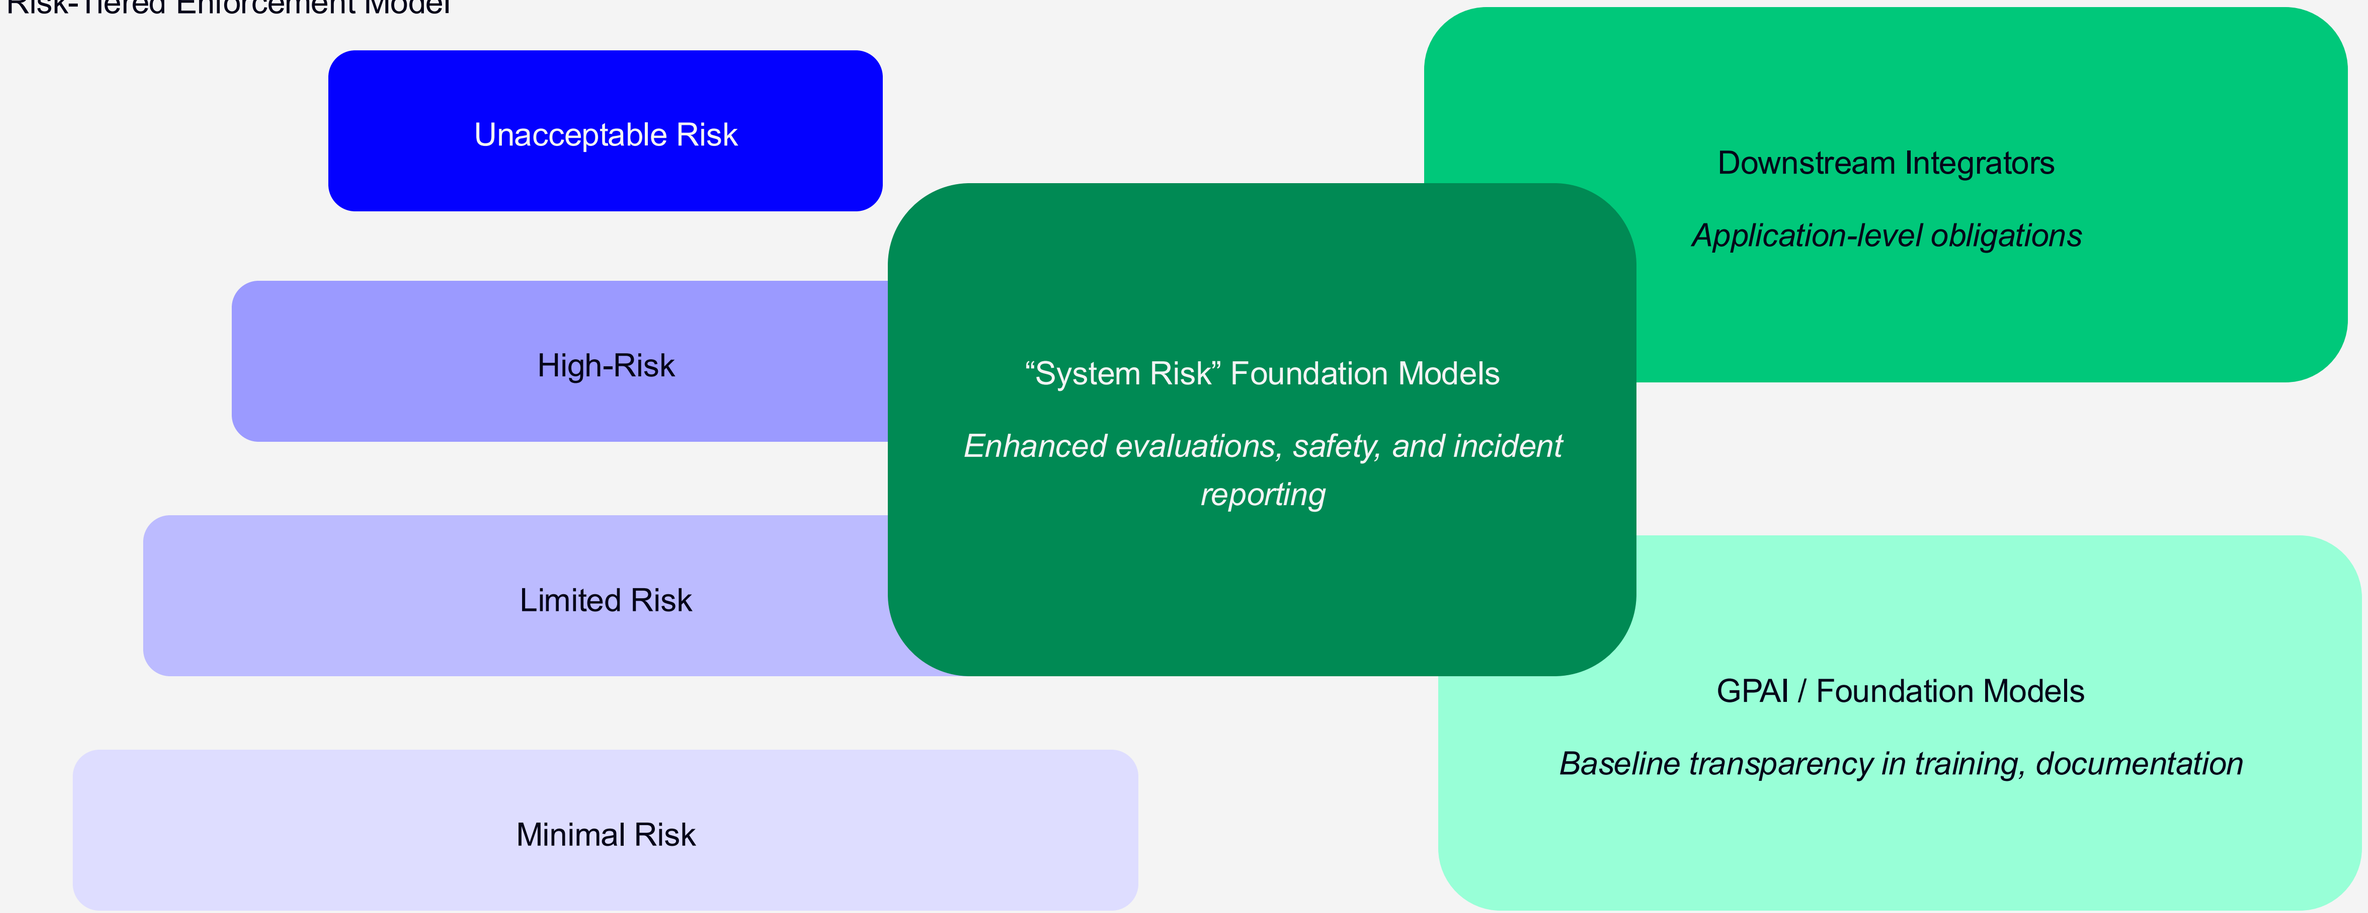

EU AI Act Application Types

	Responsibilities	Obligations
Deployer / User / Customer Downstream Integrators *Application-level obligations*	• Risk assessment in context • Human oversight • Market monitoring • Safety and usage policies also (????)	• Impact and risk assessments • ... • Usage restrictions • Incident monitoring and disclosure • ...
Model Developer GPAI / Foundation Models *Baseline transparency in training, documentation*	• Data governance • Pre-Deployment testing • Model documentation • Safety and usage policies (????)	• API terms and conditions • Usage restrictions • Training disclosures

EU AI Act

Responsibilities

Obligations

Deployer / User / Customer

...am Integrators

...evel obligations

- Risk assessment in context
- H... oversight
- ...monitoring
- ... usage policies also (????)

- Impact and risk assessments
- ...
- ...trictions
- ... and

Model De...

GPAI / Foundation Models

Baseline transparency in training, documentation

- ...ata governance
- Pre-Deployment testing
- Model documentation
- Safety and usage policies (????)

- API terms and conditions
- Usage restrictions
- Training disclosures

Other Global Standards and Guidelines

ISO / IEC 27001

Security management

Controls around data and infrastructure

Software regulation and development

NIST AI RMF

Risk identification and management

Protocols for risk management, measurement, and mitigation

Other Good Stuff

...

- ISO 23894
- IEEE
- Specific sector guidance

*These **demonstrate compliance to certain standards**, but do not guarantee compliance with local laws*

Limits of Regulation

- **1. Technology moves faster than law**
  AI systems evolve in months, while legislation and regulatory interpretation take years.
- **2. Definitions are unstable**
  It is difficult to legally define “AI,” “risk,” or “harm” in ways that remain accurate as the technology changes.
- **3. Enforcement capacity is limited**
  Regulators often lack the technical expertise and resources to audit complex models at scale.
- **4. Jurisdictional fragmentation**
  Different countries regulate AI differently, creating compliance gaps and cross-border inconsistencies.
- **5. Private actors shape outcomes**
  Companies, insurers, courts, and industry standards often influence AI governance as much as formal regulation.
- **6. Regulation cannot eliminate uncertainty**
  AI systems are probabilistic and adaptive, meaning some level of error and unintended consequence is unavoidable.

[Things Fall Apart]

DISCUSSION

EU AI Act Risk Tiers - Examples

		Definitions	Where should these go?
Risk-Tiered Enforcement Model			
Unacceptable Risk	AI systems that pose a clear threat to safety, fundamental rights, or democratic values and are therefore banned in the EU.		ofqual
High-Risk	AI systems that significantly affect people's safety or fundamental rights and are allowed only if strict compliance, documentation, and oversight requirements are met.		
Limited Risk	AI systems that pose moderate risks primarily related to deception or lack of transparency and must meet specific disclosure obligations.		
Minimal Risk	AI systems that present little to no risk to rights or safety and are largely unregulated under the Act.		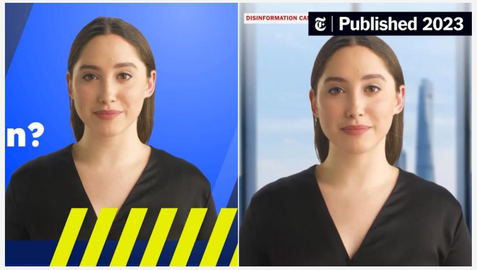

AI Litigation

So what is regulation in practice?

ofqual

Rules about implementations

- Laws (EU AI Act) require bias audit, documentation, human oversight for AI hiring systems. What a company must do prior to deployment

Guidelines on how to follow those rules

- Frameworks like NIST, ISO to shape risk assessments and bias testing requirements. “Reasonable care”

Audits and investigations

- Regulators (EEOC, EU authority) review documentation, internal systems incl. training data, validate results. Sometimes external audits

Fines, orders, recalls

- Financial penalties, suspending system for remediation. EU AI Committee, European Commission, AISI in UK

Public pressure and socializing norms

- Media, advocacy groups, customers react to bias reports; consumer pressure pushes reform in conjunction with regulatory risk

Litigation

- Applicants file discrimination law suits, which triggers discovery into model workings and governance. Court interprets compliance efforts. Everybody is sad.

Regulation is not an event, it’s an ecosystem

Case Mechanics: What is Case Law?

- Stability and predictability
  - Businesses and individuals can predict legal outcomes and plan behaviors [example – flat sharing, TV licenses]
  - Lawyers can advise clients
- Incremental adaptation
  - Building off prior precedents allows gradual refinement of rules in light of new facts and applications **(e.g., AI)**
  - Allows for deviations that must be explained for unique cases but fit into overall legal precedent structure
- Based in real applications, not hypothetical/abstract uses

Without case precedent, every case would start from scratch and be extremely reliant on particular judge biases and interpretations of the law. Case law restrains judicial power.

BUT

- Evolves slowly, behind technological application and disruption
- Only sqeakiest wheels get dealt with
- Becomes extremely complex over time

Courts don't just apply rules, ***courts create rules through precedent.***

Case Mechanics: Case Law is fundamental to our judicial system and it’s kinda the best we have

- Stability and predictability
  - Businesses and individuals can predict legal outcomes and plan behaviors [example – flat sharing, TV licenses]
  - Lawyers can advise clients
- Incremental adaptation
  - Building off prior precedents allows gradual refinement of rules in light of new facts and applications **(e.g., AI)**
  - Allows for deviations that must be explained for unique cases but fit into overall legal precedent structure
- Based in real applications, not hypothetical/abstract uses

Without case precedent, every case would start from scratch and be extremely reliant on particular judge biases and interpretations of the law. Case law restrains judicial power.

BUT

- Evolves slowly, behind technological application and disruption
- Only sqeakiest wheels get dealt with
- Becomes extremely complex over time

Courts don’t just apply rules, ***courts create rules through precedent.***

Case Mechanics: Why Case Law is Hard and Sucks

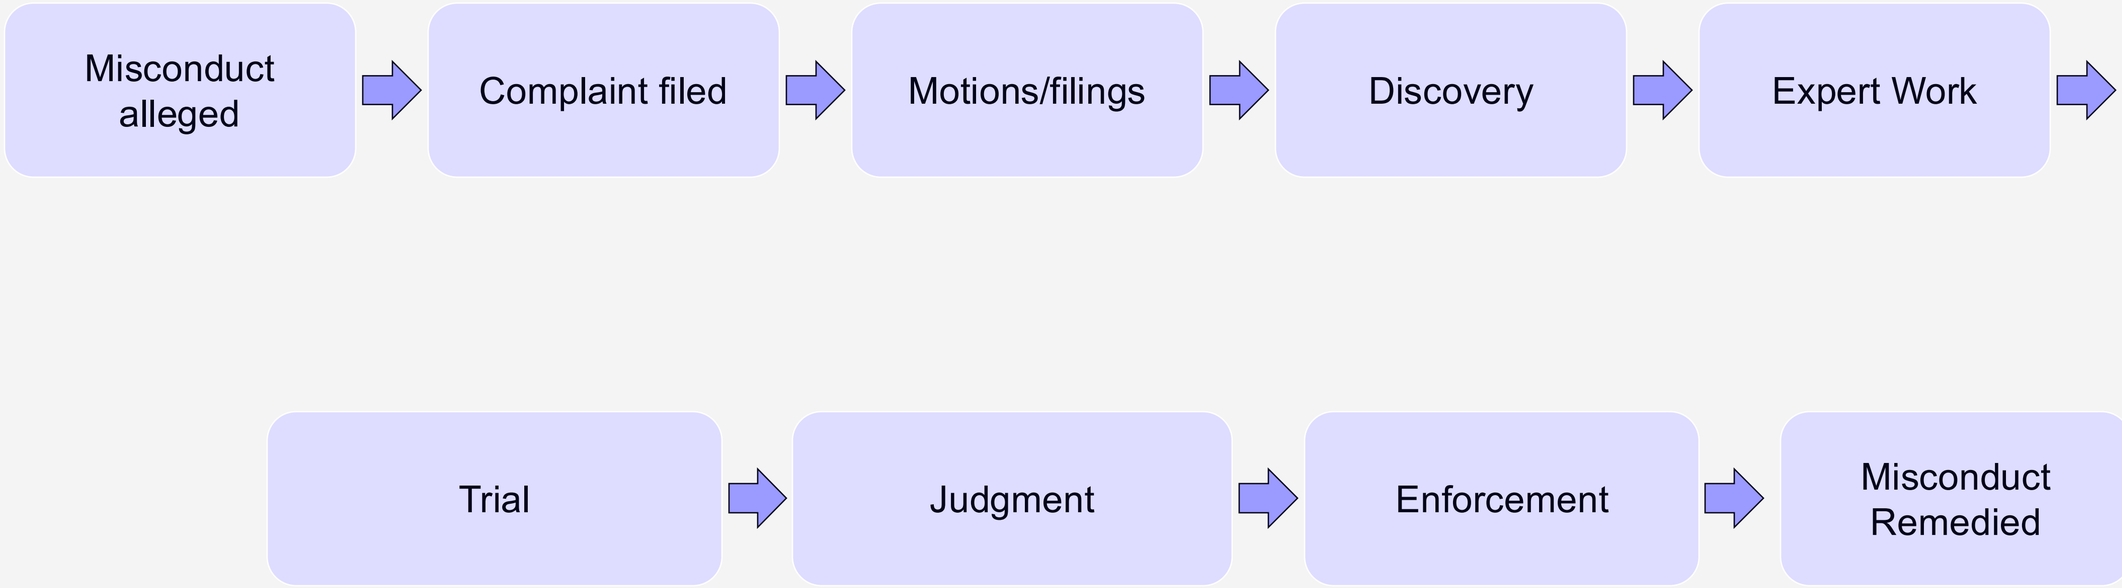

Case Mechanics: Why Case Law is Hard and Sucks

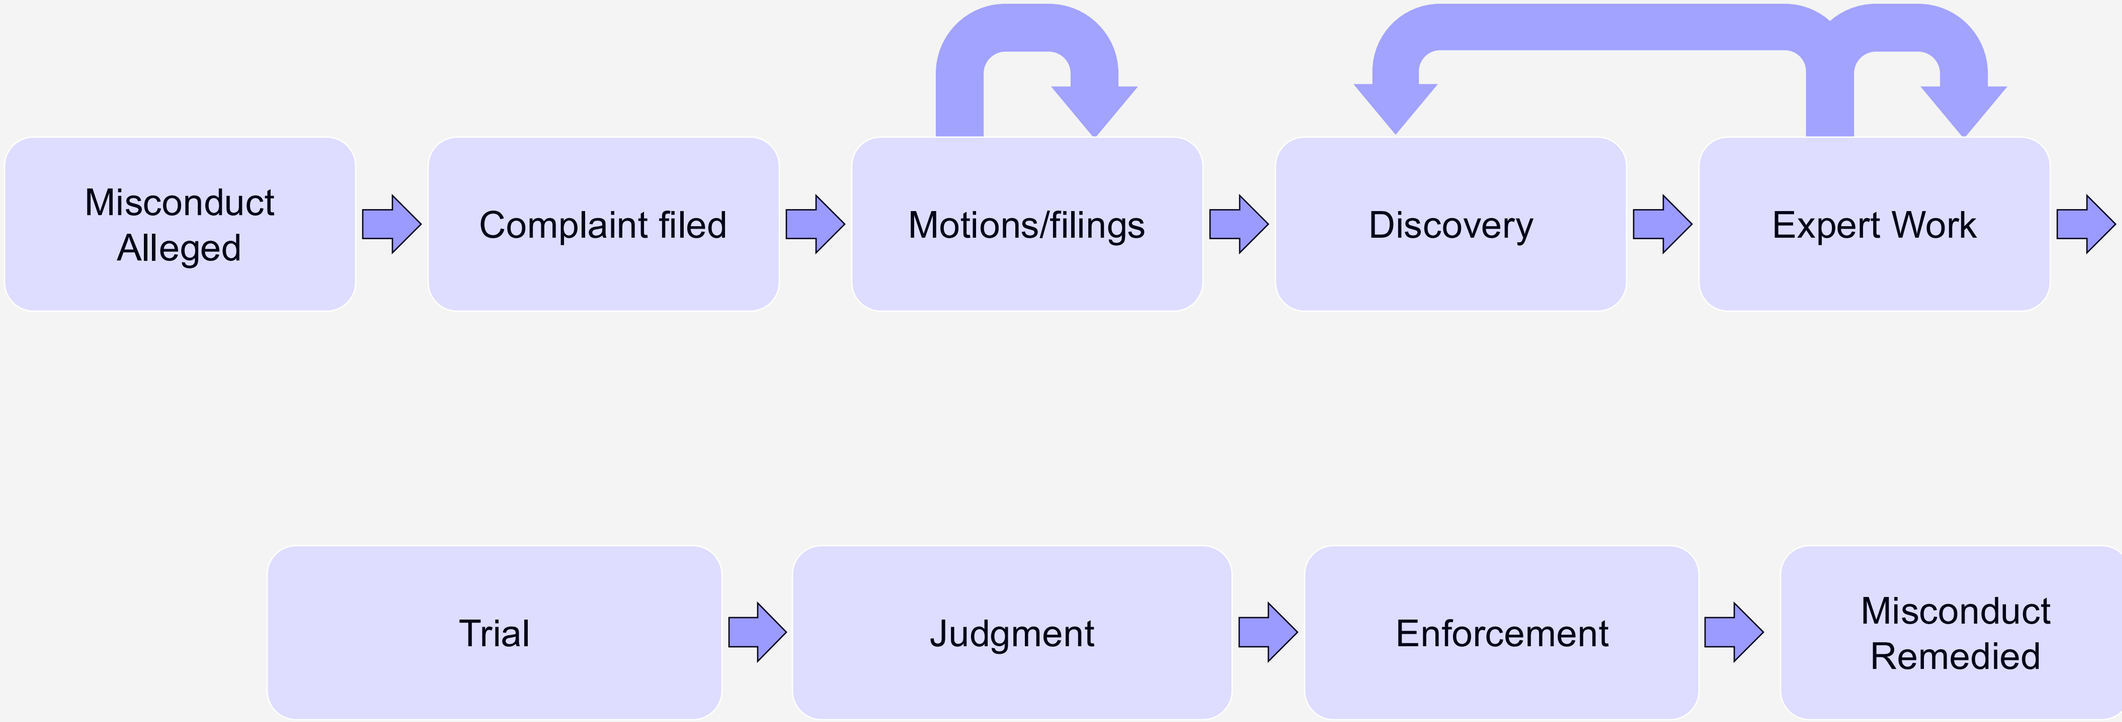

Case Mechanics: Why Case Law is Hard and Sucks

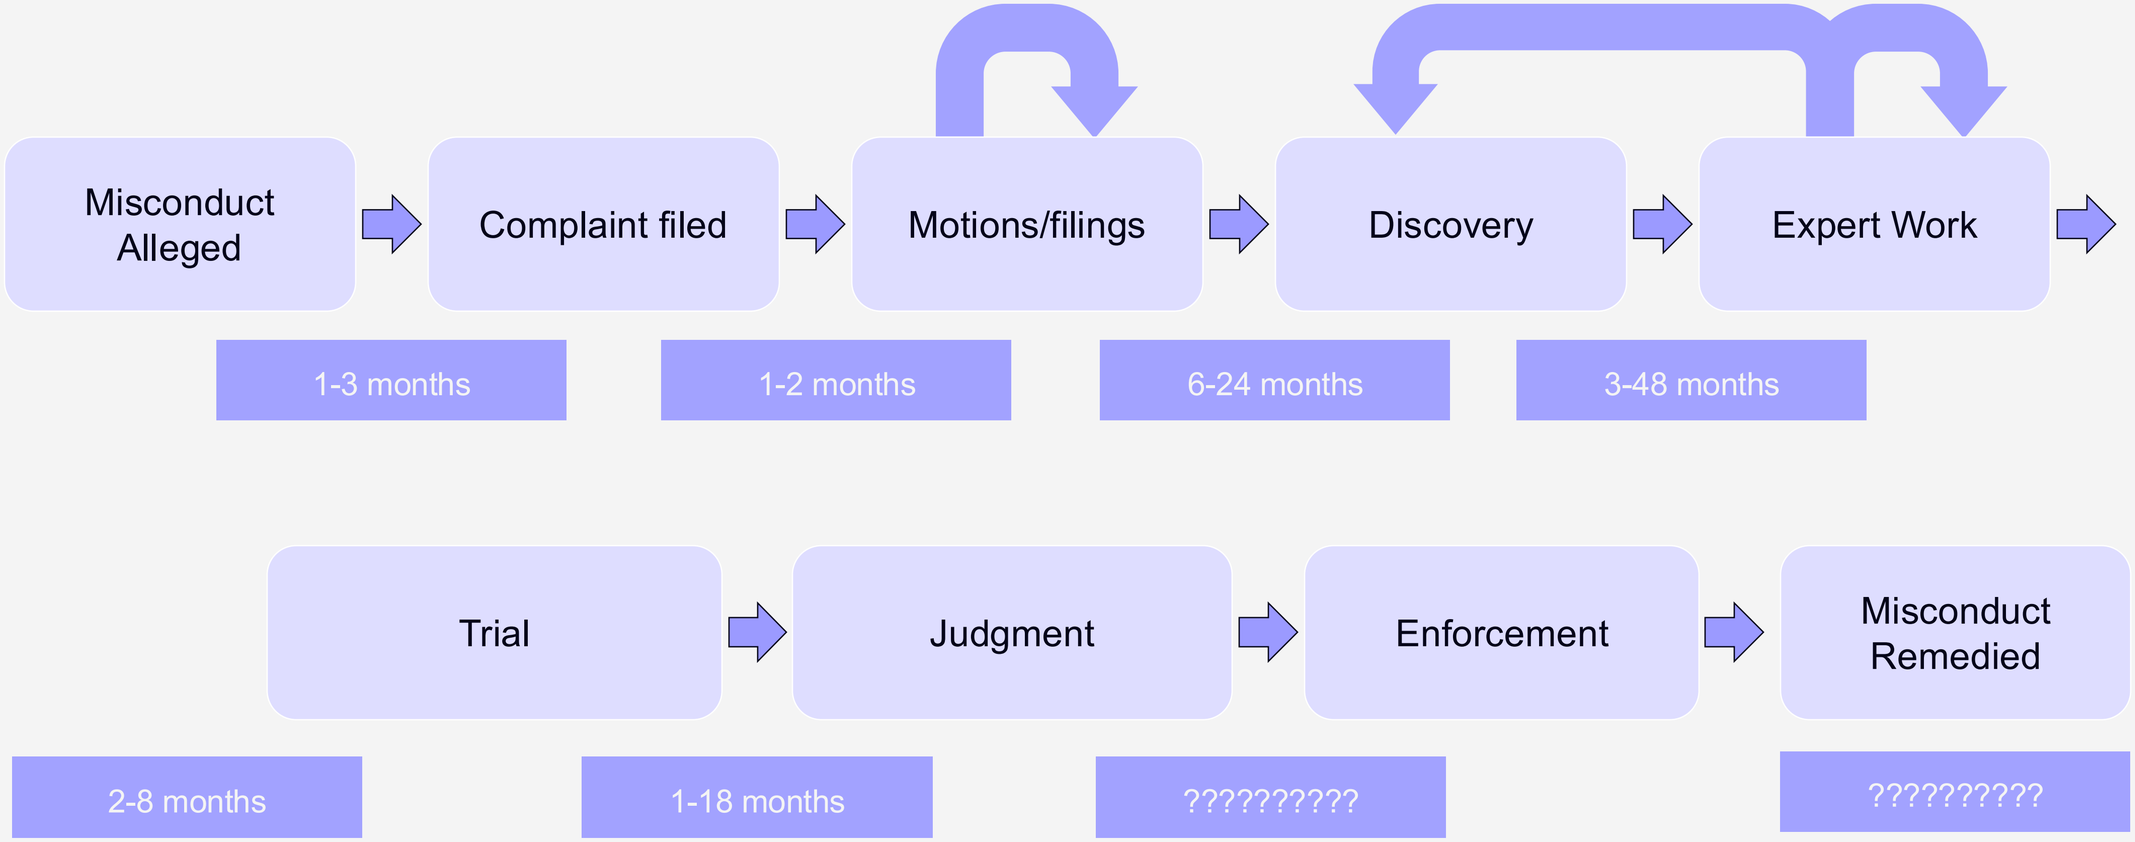

Case Mechanics: Why Case Law is Hard and Sucks

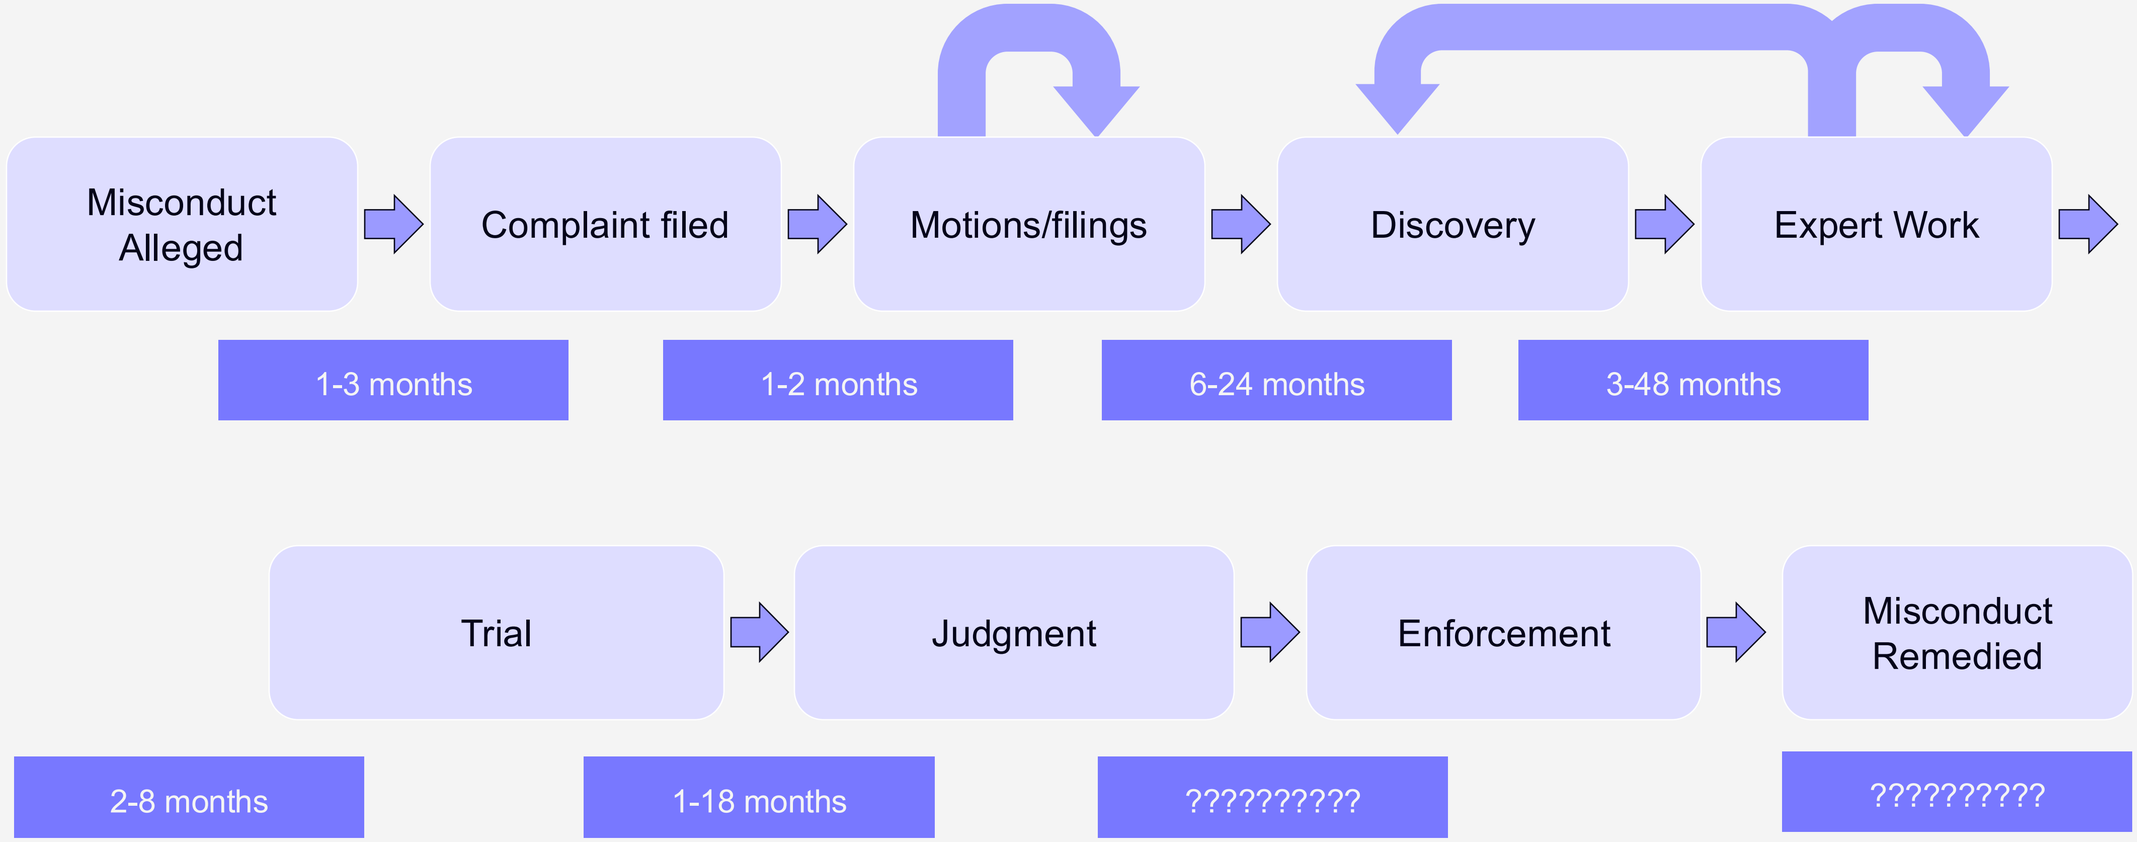

The **fastest** cases take over a year to even get to trial

Guess how complex technologies and lack of precedent impact that timeline.

Regulation vs. Litigation as Change Levers

Regulation

- Public rules
- Predictability
- *Ex ante*

Litigation

- Case-specific
- Slow
- Large burden of proof
- Large uncertainty of outcomes
- *Ex post*

- Clear obligations and precedent to motivate leadership

- Private settlements (low transparency)
- Ungeneralizable in areas without case law precedents
- In practice companies have high risk tolerance, esp. if no specific precedent

Challenges in AI Litigation

Explainability / Attribution

- Opaque “black box” generation processes do not allow for singular responsibility or authorship

Reproducibility

- Stochastic outputs make it difficult to collect reliable, repeatable data that proves harm in many circumstances

Logistics

- Data volume
- Model lineage
- Process documentation

Burden of Proof

- Lack of technical understanding make it difficult to explain and prove theories of harm to judges and juries
- Lack of case law precedent makes (some) judges extremely reticent to rule without firm understanding

Types of AI Litigation

IP / Copyright Unauthorized use of copyrighted works in training, or substantive output	**Privacy / Data Protection** Unauthorized collection, scraping, or use of personal data	**Product Liabilities** Physical or mental injury, economic loss, safety failures
Discrimination and Bias Disparate impacts in critical areas (hiring, housing, insurance, policing, health)	**Antitrust** Market dominance, exclusionary conduct, foreclosure of competition	**Misinformation** Reputational injury from defamation or hallucinations

Discussion: Types of AI Litigation Risks and Mitigations

Category	Where Did AI Go Awry?	Who Was / Is Harmed?	Who Might Be Responsible?	Why Does AI Make It Hard?
Copyright / IP				
Bias & Discrimination				
Privacy & Data Use				
Product Liability				
Antitrust				
Misinformation				

Case Studies - Discuss

1. *The New York Times v. OpenAI & Microsoft* (2023–): The New York Times sued OpenAI and Microsoft alleging that millions of its copyrighted articles were used to train generative AI models without permission and that the models can reproduce or closely summarize paywalled content, undermining its subscription business. The case centers on whether large-scale AI training constitutes copyright infringement or fair use, and whether model outputs are unlawful derivative works. The litigation is widely seen as a test case for how U.S. copyright law applies to generative AI training and output.

2. EEOC v. iTutorGroup (2023 settlement): The U.S. Equal Employment Opportunity Commission sued iTutorGroup, alleging that its AI-powered hiring software automatically rejected older applicants, programming the system to disqualify women over 55 and men over 60. The case resulted in a settlement requiring monetary relief and injunctive measures. It illustrates how existing anti-discrimination laws apply directly to automated decision systems and confirms that employers remain liable when using algorithmic screening tools.

3. Italian Data Protection Authority v. OpenAI (2023): Italy's data protection authority temporarily banned ChatGPT, alleging violations of the GDPR, including unlawful data processing and insufficient age verification safeguards. Regulators questioned the legal basis for using personal data to train the model and raised concerns about inaccurate personal information generated by the system. OpenAI restored service after implementing changes, demonstrating how privacy regulators can directly intervene in AI deployment.

4. *Mata v. Avianca* (2023): In this U.S. federal case, attorneys relied on ChatGPT to draft a legal brief that cited entirely fabricated court cases. When the inaccuracies were discovered, the court sanctioned the lawyers. While not a traditional product liability lawsuit against the AI developer, the case highlights emerging liability questions around reliance, foreseeability, and the allocation of responsibility when AI systems generate harmful or erroneous outputs.

5. European Commission Inquiry into MSFT-OpenAI Partnership (2023-2024): The European Commission examined Microsoft's multibillion-dollar investment in OpenAI and its deep integration of OpenAI models into Azure and Microsoft products to determine whether the partnership amounted to a merger or conferred anti-competitive advantages in emerging AI markets. Regulators focused on issues of control, access to compute infrastructure, preferential distribution, and whether the relationship could foreclose competitors from critical inputs or customers. The Commission concluded the partnership did not constitute a notifiable merger under EU rules at that time.

6. Case Study: *Walters v. OpenAI* (2023): A radio host in Georgia sued OpenAI for defamation after ChatGPT falsely stated that he had committed fraud and embezzlement. Although the case was ultimately dismissed, it illustrates how AI "hallucinations" can generate false statements about real individuals, raising novel questions about defamation law, fault standards, and platform responsibility when content is machine-generated rather than authored by a human.

Discussion: Types of AI Litigation Risks and Mitigations

Category	Where Did AI Go Awry?	Who Was / Is Harmed?	Who Might Be Responsible?	Why Does AI Make It Hard?
Copyright / IP	The model was trained on copyrighted books, art, music, or code without permission, or generated outputs too similar to originals.	Authors, artists, programmers, media companies (economic harm).	Model developer, company deploying the system, possibly users in limited cases.	Training requires massive datasets scraped at scale. Models don't store exact copies (usually), so it's unclear whether "learning from" content is the same as copying it.
Bias & Discrimination	The AI system made systematically unfair decisions (e.g., in hiring, lending, admissions, insurance).	Job applicants, borrowers, tenants, patients — often members of protected groups.	The company using the AI system, the vendor that built it, sometimes both.	Bias can emerge from training data even without intent. Models are opaque, making it hard to prove why a decision happened. Responsibility may be split between builder and deployer.
Privacy & Data Use	Personal data was scraped, stored, inferred, or reused in unexpected ways (e.g., facial recognition, voice cloning, data scraping).	Individuals whose personal data was collected or inferred.	Data collectors, AI developers, companies deploying the system.	"Public" data can still feel private. AI can infer new sensitive information. Once data is embedded in training, it's difficult to remove.
Product Liability	An AI-enabled system malfunctioned or made a bad decision that caused injury or financial loss (e.g., self-driving crash, medical misdiagnosis).	Users, bystanders, customers, patients.	Software developer, hardware manufacturer, system operator, integrator.	AI systems are probabilistic and adaptive. It's difficult to trace a specific decision back to a specific design flaw. Traditional product law assumes static, predictable products.
Antitrust	A few companies gained dominant control over key AI resources (compute, data, models), limiting competition.	Competitors, startups, consumers (less innovation, higher prices, reduced choice).	Large AI companies, cloud providers, vertically integrated tech firms.	AI requires enormous compute and data advantages. Network effects make it hard for new entrants. It's unclear when scale becomes illegal dominance.
Misinformation	The AI generated false statements about real people (e.g., inventing crimes, misconduct, fake credentials)	Individual reputations were harmed	User misuse, hosting platform, possibly model developer	Systems generate text probabilistically – no human "author" to have invented statement. Hard to determine who "said" something

Takeaways

If nothing else please remember:

“

Regulation Facilitates Innovation

That’s it.

“

AI Litigation is a Complex and Evolving (Mine) Field

Avoid it.

“

Governance and Process now Reduces Risk and Litigation Later

For real, avoid litigation.

Practical ways to use regulation and legislation to bully your company into doing the right thing

1

Frame everything as a business concern

- Businesses vary in how much they care about “doing the right thing” or “following the law” – but they take *risk* seriously.
- Following regulations is frequently healthiest for business growth in the long term

2

Know the regulation that *could* apply to your situation, and the responsibilities of your role in the organization

- Design with risk in mind, literally always.
  - ALL AI applications fall under some sort of regulation in the UK / US / EU; **engineers need to anticipate foreseeable harm.**

3

Build auditability and documentation into engineering systems and decisions

- In legal situations, the company that can’t explain what they’ve done, how, or why, is the company that looks irresponsible and loses.
- When you come to the company/ regulators with concerns, you need data.

4

Participate in governance and policy discussions in your organization

- Engineering-forward orgs, SaaS, finance, healthcare, education, any company that implements or uses AI desperately need people with technical expertise to understand the big picture

Q & A