THE BULLETEEN
APR '26

DISORDER

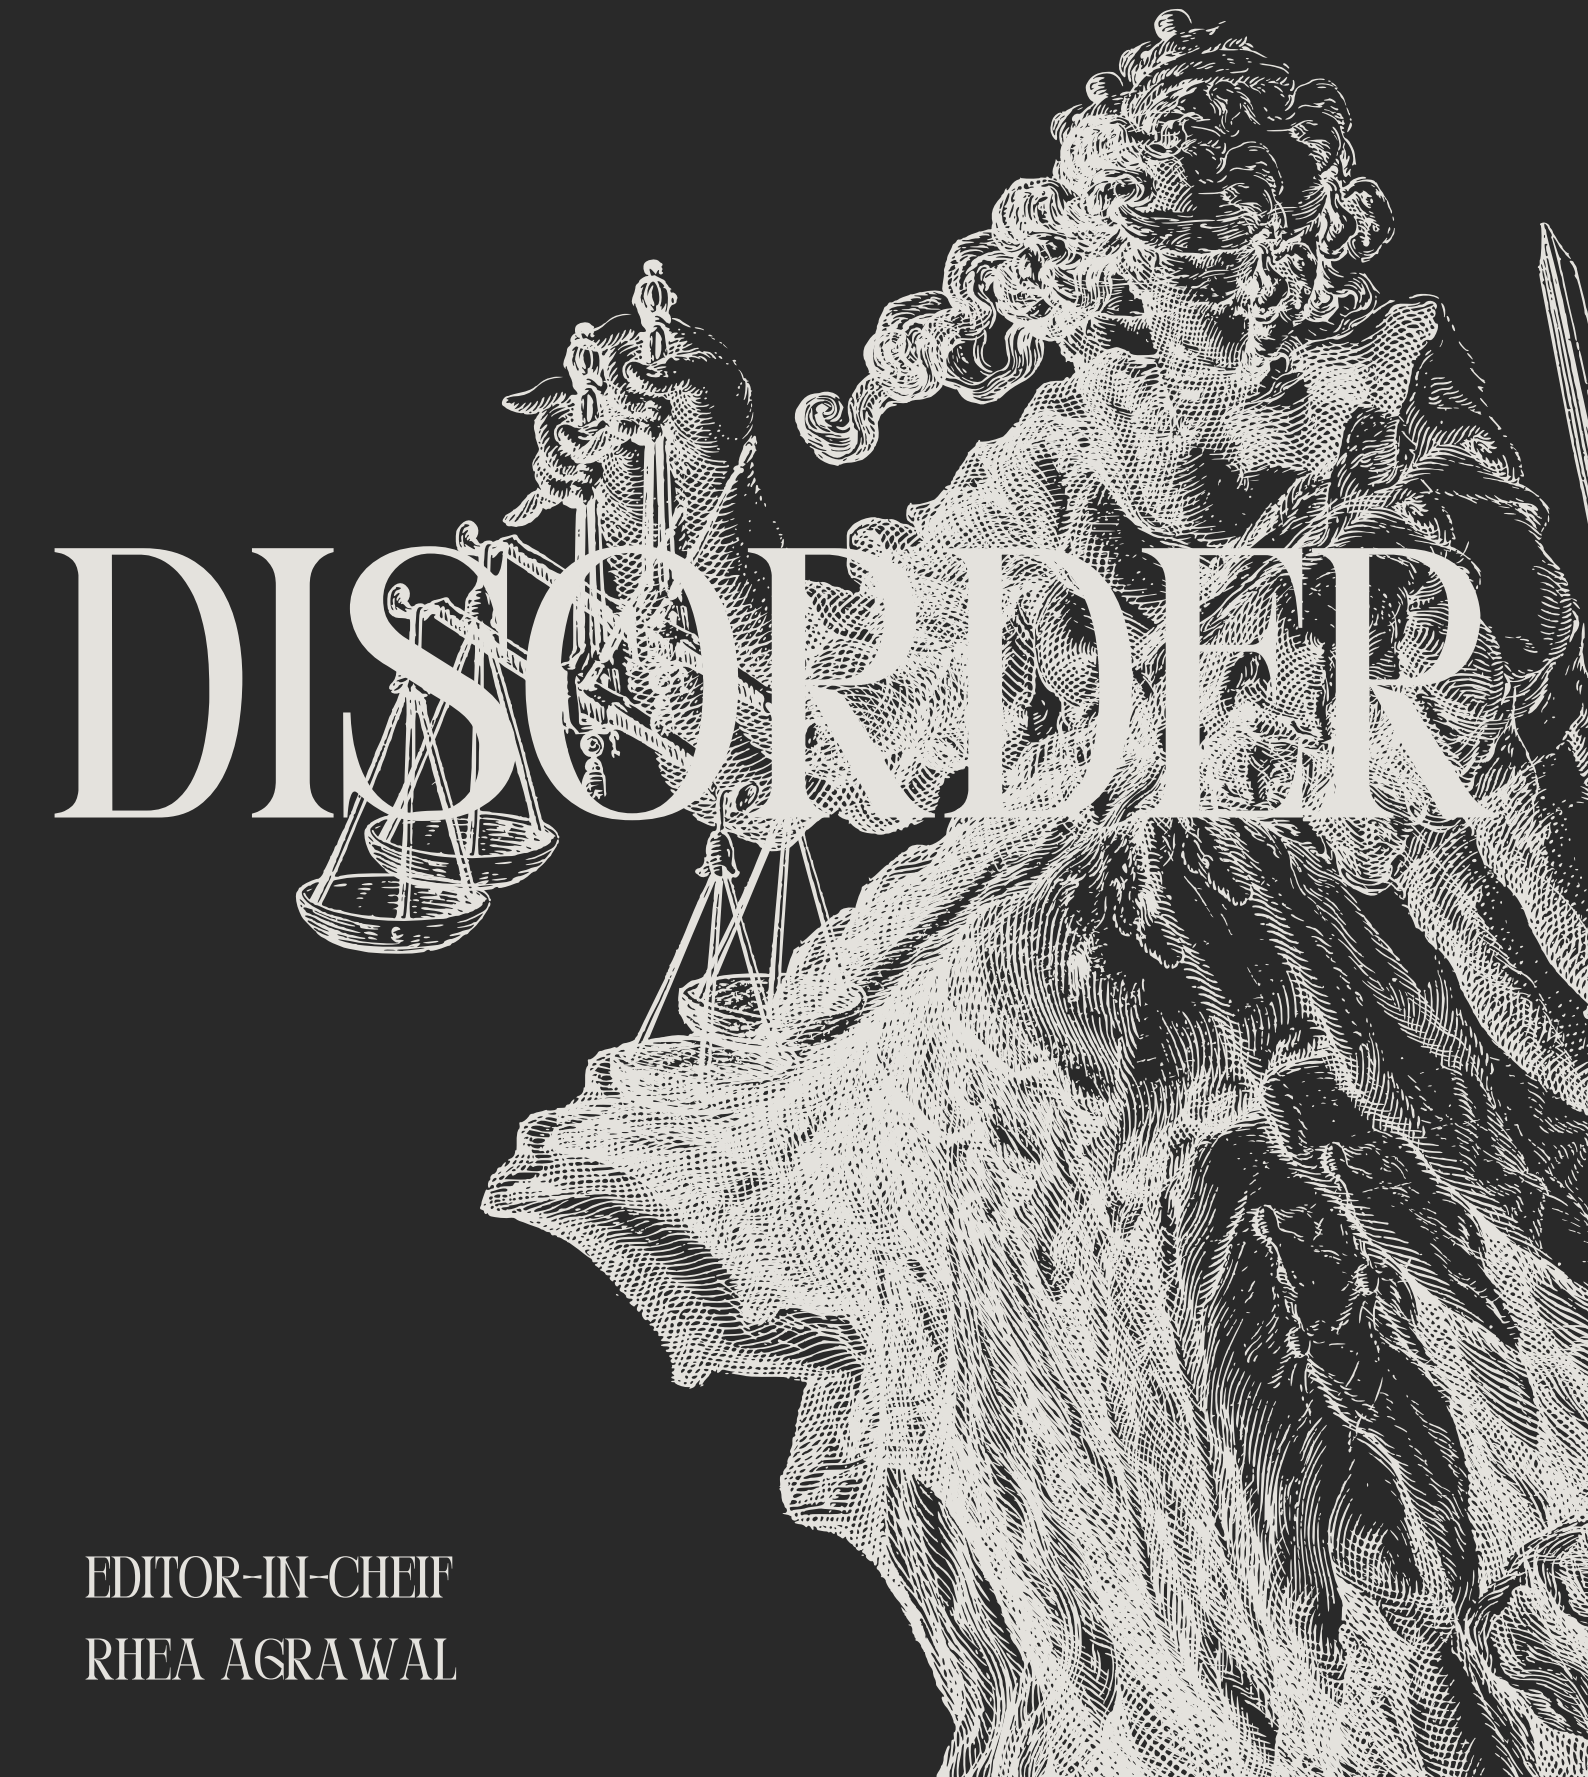

EDITOR-IN-CHEIF
RHEA AGRAWAL

TABLE OF CONTENTS

THE DNA DILEMMA

Few technologies have transformed the courtroom as dramatically as DNA analysis. What began as a laboratory curiosity in the mid-1980s has become the gold standard of forensic evidence, capable of exonerating the innocent and convicting the guilty with a degree of certainty that was unimaginable to earlier generations of lawyers and judges. Yet as the science grows more powerful, it raises profound legal and ethical questions that the justice system is only beginning to grapple with.

The story of DNA in law enforcement begins with British geneticist Alec Jeffreys, who in 1984 discovered that certain regions of the human genome vary dramatically between individuals. By 1986, this technique was being used in a murder investigation in Leicestershire, England, ultimately both exonerating an innocent suspect and identifying the true perpetrator. The age of forensic genomics had arrived.

Today, law enforcement agencies around the world rely on DNA databases to solve crimes. In the United States, the FBI's Combined DNA Index System (CODIS) contains profiles from tens of millions of individuals, linking crime scenes to suspects with stunning efficiency. A single hair follicle, a drop of saliva, or skin cells left on a doorknob can be enough to place a person at the scene of a crime.

But the science has grown far more sophisticated, and far more contested, in recent years. Probabilistic genotyping software now allows analysts to interpret complex DNA mixtures involving multiple contributors, using statistical algorithms to calculate the likelihood that a particular individual's genetic material is present. Critics argue that these proprietary "black box" programs are difficult for defense attorneys to scrutinize, raising serious due-process concerns. When defendants cannot meaningfully challenge the software that helped convict them, the fundamental legal right to confront evidence against oneself is undermined.

Even more provocative is the rise of forensic DNA phenotyping, which attempts to predict a suspect's physical appearance, eye color, hair color, skin tone, even facial structure, from genetic material alone. While the technology holds promise for generating investigative leads when no database match exists, it also carries a troubling risk of racial profiling. If an algorithm predicts that a suspect is likely of a particular ethnic background, investigators may unconsciously narrow their focus in ways that embed bias into the process.

The use of familial DNA searching adds another layer of complexity. By identifying partial matches in a database, investigators can sometimes locate a suspect's close relatives, then work outward to identify the suspect themselves. This technique has solved high-profile cold cases, including the 2018 identification of the Golden State Killer. But it also means that anyone with a relative in a DNA database has, in effect, partial genetic representation in that database, whether they consented to it or not. The legal frameworks governing this kind of search vary widely and remain deeply unsettled.

Then there is the question of wrongful conviction. The Innocence Project, which uses DNA testing to review cases of potential wrongful conviction, has helped exonerate more than 375 people in the United States since 1992. Many of these individuals spent decades in prison for crimes they did not commit, convicted in part on forensic evidence that was later shown to be flawed or misinterpreted. The lesson is humbling: even scientifically grounded evidence can be misused, misunderstood, or overstated in the courtroom.

Training is a critical bottleneck. Forensic scientists require rigorous education in genetics, statistics,

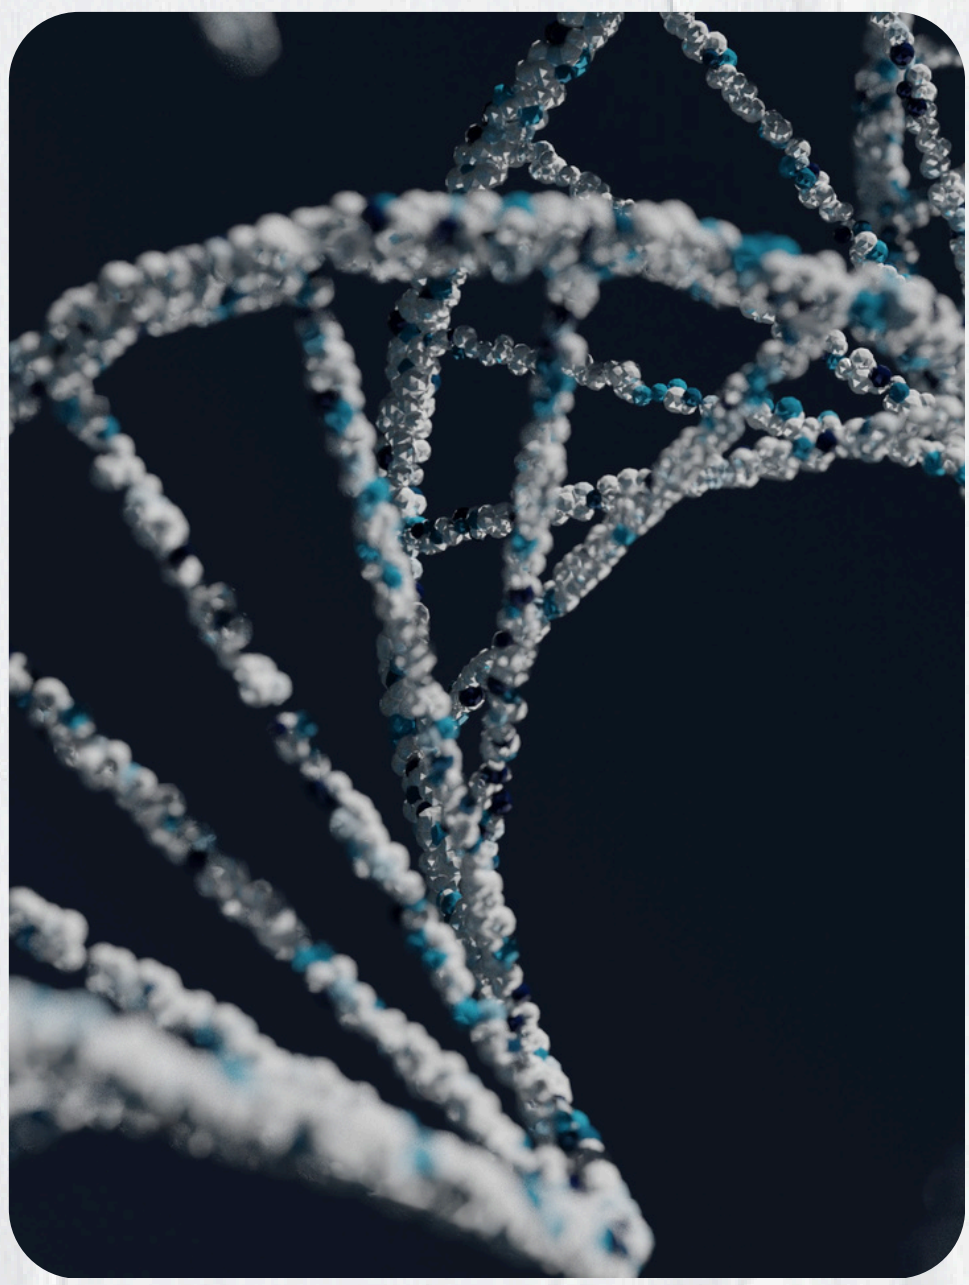

and laboratory procedure, yet expert witnesses are not always well-equipped to communicate probabilistic reasoning to juries. When an analyst says that the odds of a random match are "one in a trillion," jurors may accept this figure uncritically, not understanding the assumptions built into that calculation or the contexts in which it might break down.

Legal scholars and scientists are increasingly calling for reform. Proposals include mandatory disclosure of the source code of forensic software, standardized training for expert witnesses, independent auditing of crime laboratories, and clearer rules about when probabilistic evidence may be presented in court. Some jurisdictions are moving toward these reforms; others lag far behind.

The intersection of genomics and justice will only grow more complex as the technology advances. Whole-genome sequencing, once prohibitively expensive, is becoming cheaper every year. The day is approaching when a complete genetic profile can be obtained from trace evidence in hours, at minimal cost. This will bring new investigative power, and new legal responsibilities. Society must decide not just what the science can do, but what it should do, and under what conditions, constraints, and oversight. The courtroom has always been a place where truth is contested. DNA has not changed that, it has simply raised the stakes.

ALGORITHMS IN THE DOCK

In courtrooms across the United States, judges are increasingly turning to artificial intelligence tools to help them make one of the most consequential decisions in the legal system: how long a convicted person should spend in prison. Risk assessment algorithms, with names like COMPAS, PSA, and LSI-R, promise to bring objectivity and consistency to a process long criticized for its reliance on individual judicial discretion. The reality, however, is considerably more complicated, and considerably more troubling.

These tools work by collecting data about a defendant, criminal history, age, employment status, education level, family background, and running it through a statistical model trained on historical data to generate a score.

A high score indicates a higher predicted likelihood of reoffending; a low score suggests the opposite. Judges are then invited to consider this score alongside other factors when deciding on bail, sentencing, and parole.

Proponents argue that algorithmic tools can reduce the inconsistency and implicit bias that plague human decision-making. Research has long shown that sentences for identical crimes can vary dramatically depending on the judge, the day of the week, and whether the judge has recently had lunch. If a data-driven tool can smooth out these arbitrary disparities, the argument goes, justice becomes more equitable.

But a landmark 2016 investigation by ProPublica exposed a critical flaw in this reasoning. Analyzing the COMPAS tool used in Florida, journalists found that Black defendants were nearly twice as likely as white defendants to be falsely flagged as future criminals, while white defendants were more likely to be incorrectly labeled low-risk. The algorithm, it turned out, was reproducing and amplifying the racial disparities already embedded in the criminal justice system, because it was trained on historical data that reflected those very disparities.

This is the core paradox of algorithmic fairness: a model trained on biased data will produce biased predictions, regardless of how sophisticated its mathematics. Arrest rates, conviction rates, and incarceration rates are not neutral measures of criminal behavior, they are measures of how the justice system has historically treated different populations. When an algorithm treats these as proxies for future risk, it codifies past injustice into future decisions.

The due-process issue is profound. In a system built on the right to confront evidence and challenge the basis for one's punishment, a "black box" algorithm presents a fundamental obstacle. If a company's intellectual property claims prevent a defendant from understanding why they received a high-risk score, the adversarial process that underlies fair adjudication is compromised.

Beyond bias and transparency, there is a deeper philosophical question: should future behavior, which has not yet occurred, ever be the basis for present punishment? Traditional retributive justice focuses on what a person has done, not on statistical predictions about what they might do. Risk-based sentencing, by contrast, moves toward a model of preventive detention that raises serious civil liberties concerns. Punishing people for crimes they have not yet committed, on the basis of group-level statistics, sits uneasily with foundational principles of individual rights.

The challenge is not simply technical. It requires society to make explicit value choices about what fairness means, whose past data is used to predict whose future, and whether efficiency can ever justify inequality. Technology can process data at superhuman speed, but it cannot tell us what justice requires. That judgment remains irreducibly human — and must remain subject to human accountability. As AI tools become more powerful and more pervasive in the legal system, the imperative to govern them wisely has never been greater.

CYBERCRIME AND JURISDICTION: WHEN DIGITAL BORDERS CHALLENGE LEGAL BOUNDARIES

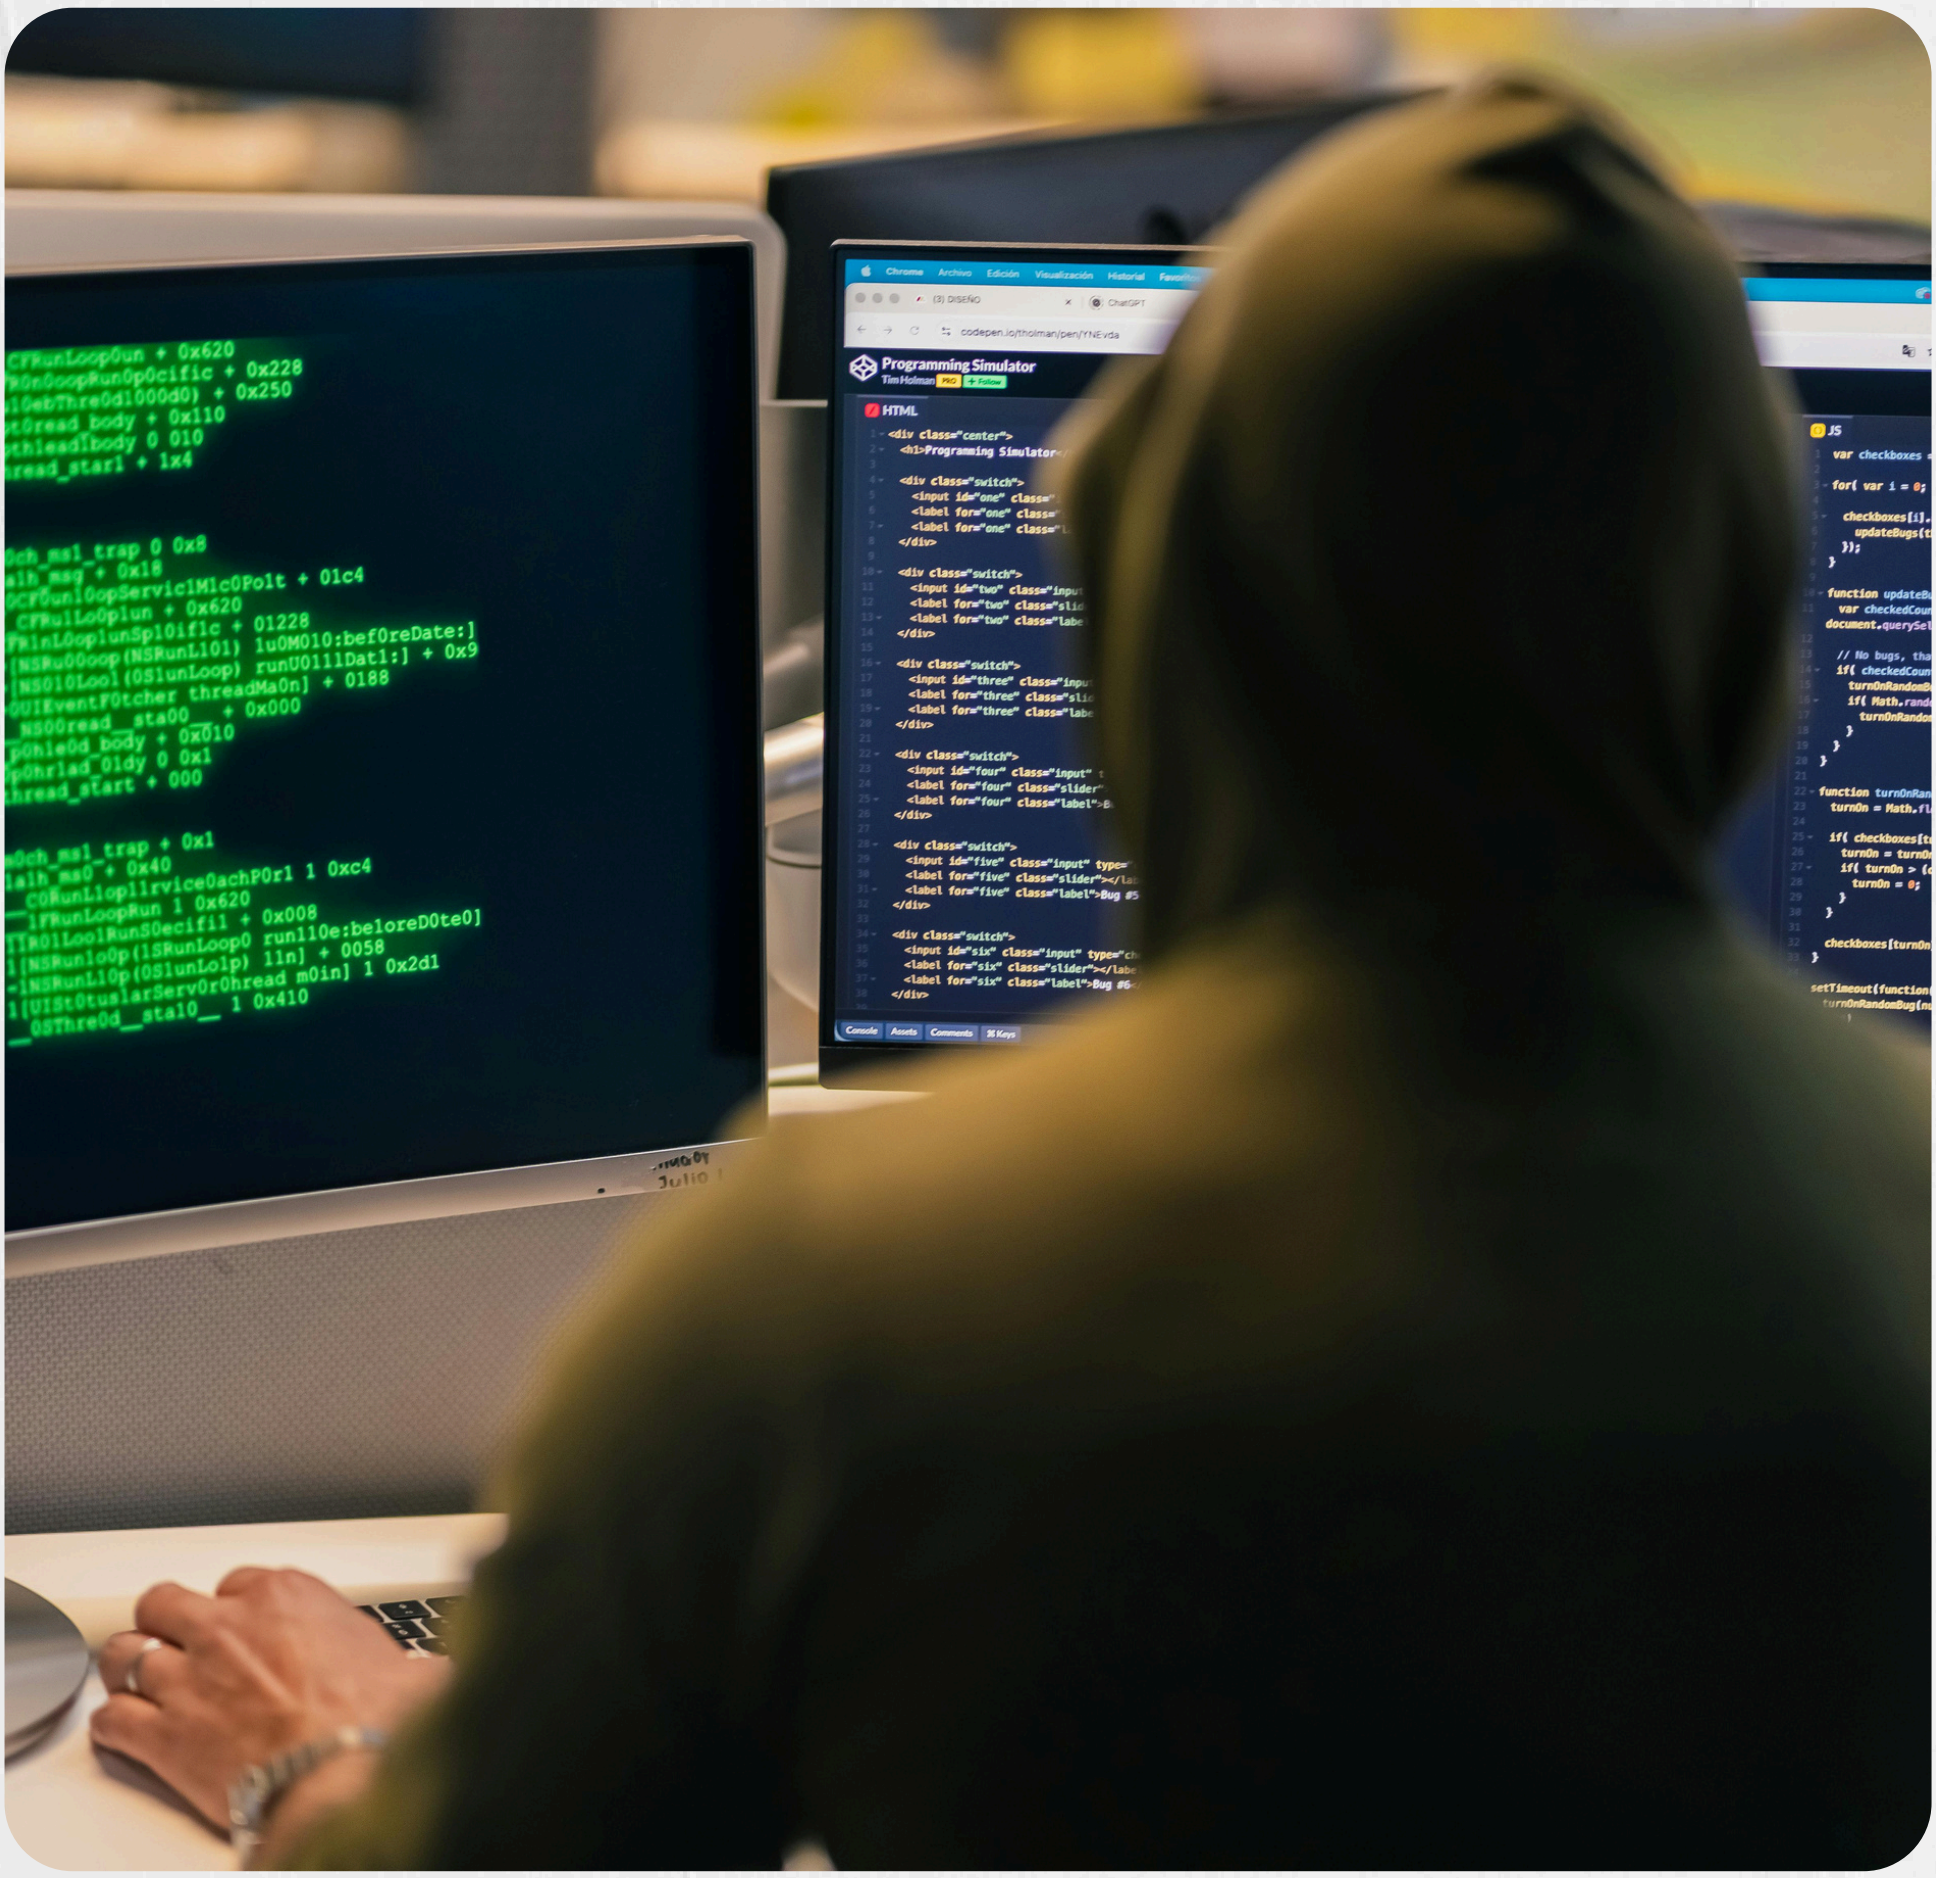

On a Tuesday morning in Bucharest, a programmer writes code. By Wednesday, that code has infected hospital computers in California, locked up patient records in Manchester, and extorted payments in bitcoin to an account registered in the Seychelles. By Thursday, the programmer is back at his keyboard, working on the next attack. By the time any law enforcement agency has identified what happened, traced the malware, and begun the process of seeking international cooperation, months may have passed. The perpetrator may have moved, changed identities, or been shielded by a government with no interest in extraditing him.

This is the fundamental dilemma of cybercrime: the technology moves faster than the law, and the law is constrained by borders that the technology simply ignores.

Cybercrime is among the fastest-growing categories of criminal activity in the world. Ransomware attacks, phishing scams, corporate espionage, election interference, child exploitation networks, and financial fraud collectively cost the global economy hundreds of billions of dollars each year. Yet prosecution rates remain stubbornly low. The reasons are both technical and legal, and the interplay between the two defines the central challenge of digital law enforcement.

From a technical standpoint, cybercriminals have numerous tools for concealing their identities and locations. Virtual private networks (VPNs), the Tor anonymization network, encrypted messaging applications, and cryptocurrency payment systems make it genuinely difficult to attribute an attack to a specific individual. Even when investigators do identify a suspect, the digital evidence trail, logs, IP addresses, metadata, must be collected and preserved in ways that meet legal standards for admissibility, which vary dramatically from country to country.

From a legal standpoint, criminal law has historically been territorial. A country's laws apply within its borders; beyond them, it must rely on treaties and the cooperation of foreign governments. The principal international instrument for cybercrime law is the Budapest Convention on Cybercrime, adopted by the Council of Europe in 2001. Now ratified by 68 countries, the convention establishes common definitions for cybercrime offenses and creates mechanisms for cross-border evidence sharing and mutual legal assistance.

But the Budapest Convention has significant limitations. Major cybercrime source countries, including Russia, China, Brazil, and India, have not ratified it. Russia in particular has long been accused of tolerating or actively facilitating cybercriminal activity directed at Western targets, providing de facto safe havens for criminal networks. Without meaningful cooperation from these countries, international law enforcement efforts are severely hampered.

Traditional mutual legal assistance processes, designed for an era of paper records and physical evidence, can take months or years to yield results. By the time a formal request for evidence is processed through diplomatic channels, the data may be long gone.

In response, some jurisdictions have adopted more aggressive unilateral approaches. American law enforcement, empowered by statutes like the Computer Fraud and Abuse Act and the CLOUD Act, has sometimes asserted extraterritorial jurisdiction over cybercrime cases, pursuing prosecutions of foreign nationals for attacks on American infrastructure. While effective in high-profile cases, this approach can create diplomatic friction and raises questions about sovereignty and reciprocity.

The private sector plays an increasingly important role. Major technology companies, Microsoft, Google, Meta, and others, maintain their own cybersecurity teams and threat intelligence units that often operate faster and more nimbly than government agencies. Microsoft's Digital Crimes Unit has disrupted numerous botnet operations and nation-state hacking groups through civil litigation and technical countermeasures. But private enforcement raises its own questions about accountability and the appropriate boundaries between corporate and governmental power.

Emerging legal frameworks are attempting to close the gaps. The European Union's General Data Protection Regulation has established new standards for data sovereignty that have implications for cybercrime investigation. A new UN treaty on cybercrime is currently under negotiation, though its prospects are uncertain and its scope contested.

The underlying challenge is structural: building legal frameworks that are nimble enough to keep pace with rapidly evolving technology, global enough to reach across borders, and principled enough to protect civil liberties while enabling effective law enforcement. There are no easy answers, but the stakes, for economic security, critical infrastructure, and democratic institutions, could not be higher.

THE SURVEILLANCE STATE

In cities around the world, cameras are watching. In London, an estimated 500,000 closed-circuit television cameras monitor public spaces. In Shenzhen, China, a network of millions of cameras is integrated with facial recognition systems capable of identifying jaywalkers and displaying their faces on public screens. In Chicago, police have access to a network of sensors, cameras, and license plate readers linked by predictive analytics software. The age of pervasive surveillance has arrived, and with it, a profound legal and ethical reckoning.

The technology enabling modern surveillance has advanced at extraordinary speed. High-resolution cameras have become cheap and ubiquitous. Facial recognition algorithms, trained on vast datasets of human faces, can now match an individual against a database of millions of images in seconds. Gait recognition systems can identify a person by the way they walk, even when their face is obscured. Location data from smartphones allows detailed reconstruction of an individual's movements over time. Taken together, these technologies create the possibility of a surveillance architecture of unprecedented scope and power.

Law enforcement agencies argue that this technology saves lives. CCTV footage has been instrumental in solving terrorist attacks, murders, and kidnappings. Facial recognition has helped identify suspects in cases where no other leads existed. Automatic license plate readers have recovered stolen vehicles and located fugitives. Predictive policing systems, which analyse historical crime data to forecast where crimes are likely to occur, are credited with reducing crime rates in some cities.

But the legal and civil liberties community has raised serious and well-documented concerns. Facial recognition technology, in particular, has been shown to perform significantly worse on darker-skinned faces and on women, a problem documented in landmark research by MIT's Joy Buolamwini. The implications for law enforcement are grave: if a facial recognition system misidentifies a person as a suspect, that person may face arrest, detention, and prosecution for a crime they did not commit. Several such cases have already occurred in the United States, all involving Black men falsely identified by automated systems

The legal framework governing police surveillance in most democracies was written for a different era. The Fourth Amendment to the United States Constitution protects against unreasonable searches and seizures, but courts have long struggled to apply this principle to digital surveillance. The Supreme Court's 2018 decision in Carpenter v. United States held that police must obtain a warrant before accessing historical cell phone location data — a significant victory for digital privacy. But vast swaths of surveillance activity remain unregulated or under-regulated, operating in legal gray zones that courts and legislatures have not yet addressed.

The question of consent is particularly thorny. When a person walks down a public street, do they implicitly consent to being recorded, recognized, and tracked? Traditional legal doctrine holds that there is no reasonable expectation of privacy in public spaces. But this doctrine was developed in an era when surveillance was expensive, limited, and temporary. The pervasive, permanent, and automated surveillance now technically possible is qualitatively different from a police officer observing a public square, and the law has been slow to recognize this distinction.

Some jurisdictions are beginning to push back. San Francisco became the first major American city to ban city agencies from using facial recognition technology in 2019, followed by several other cities and states. The European Union's Artificial Intelligence Act, adopted in 2024, places strict limits on the use of real-time facial recognition in public spaces by law enforcement. These regulatory responses reflect a growing consensus that the deployment of powerful surveillance tools must be subject to democratic oversight and meaningful legal constraint.

Privacy-preserving technologies, such as differential privacy and federated learning, may allow certain analytical goals to be achieved without centralizing sensitive personal data. Algorithmic auditing frameworks can evaluate surveillance systems for bias and accuracy before deployment. Transparency requirements can ensure that the public knows what tools are being used in their name.

But technology alone cannot resolve the fundamental value conflict at the heart of surveillance: the tension between collective security and individual liberty. That tension must be navigated through democratic deliberation, legal accountability, and a commitment to the principle that power exercised in the name of public safety must itself be subject to the rule of law.

Pexels. https://www.pexels.com/photo/white-hospital-beds-236380/

INTELLECTUAL PROPERTY IN THE AGE OF AI

When an artificial intelligence system generates a novel, composes a symphony, designs a drug molecule, or produces a photorealistic painting, a deceptively simple question arises: who owns it? The answer matters enormously, for creative industries, for scientific innovation, for the technology companies building AI systems, and for the human artists and inventors whose work trained those systems. And right now, the law has no satisfying answer.

Intellectual property law, encompassing copyright, patents, trademarks, and trade secrets, was built around a foundational assumption: that creativity and invention are distinctively human activities. Copyright protects original works of authorship. Patents protect novel inventions. In both cases, the law has historically required a human author or inventor to hold rights. Machines, in this framework, are tools, like a paintbrush or a typewriter, and the person wielding them holds any resulting rights.

Artificial intelligence challenges this assumption at every level. Modern generative AI systems, large language models, image generators, music synthesis tools, are not simply executing programmed instructions. They are generating outputs through statistical pattern recognition across enormous training datasets, producing results that are often genuinely novel and, by many aesthetic or functional measures, impressive. When a system like DALL-E produces an image that has never existed before, or when a protein structure prediction model like AlphaFold identifies a previously unknown molecular configuration, the traditional toolkit of intellectual property law begins to creak.

The copyright question came into sharp focus with a 2023 decision by the United States Copyright Office, which refused to register copyright in an AI-generated artwork created autonomously by a system called DABUS. The Office's position was clear: copyright requires human authorship, and a work produced without any human creative input is not eligible for protection. A federal court affirmed this position. The practical implication is significant: AI-generated works, absent meaningful human creative contribution, fall into the public domain the moment they are created.

The patent question is equally contested. Stephen Thaler, the creator of DABUS, also attempted to name the AI system as the inventor on patent applications in multiple countries. Most jurisdictions rejected these applications on the ground that inventors must be natural persons. The United Kingdom's Supreme Court reached the same conclusion in 2023. The emerging consensus is that an AI cannot hold intellectual property rights, but the question of who, if anyone, should hold such rights in AI-generated inventions remains open.

Meanwhile, a parallel legal battle is raging over the training data used to build AI systems. Generative AI models are trained on vast corpora of text, images, and other creative works scraped from the internet, often without the permission of the original creators.

A wave of lawsuits has been filed against AI companies by authors, visual artists, musicians, and news organizations, arguing that this training constitutes copyright infringement on a massive scale.

Courts must grapple with questions about whether training an AI on copyrighted material constitutes copying in a legally relevant sense, whether the outputs of AI systems are substantially similar to the works they were trained on, and whether the fair use doctrine can shield AI training from liability. Early court decisions have produced mixed results, and the issue is likely to reach appellate courts and eventually the Supreme Court.

International dimensions complicate matters further. Intellectual property law varies significantly across jurisdictions. Japan has adopted a relatively permissive approach to AI training data, while the European Union's AI Act and Copyright Directive take more restrictive positions. The result is a patchwork of conflicting rules that creates uncertainty for global technology companies and for the international creators whose work flows across borders.

The fundamental challenge is that intellectual property law was designed to incentivize human creativity and invention by granting temporary monopolies in their outputs. AI changes the problem in ways that may require not just updated rules, but a reconceptualization of what intellectual property law is for, and who it is meant to serve.

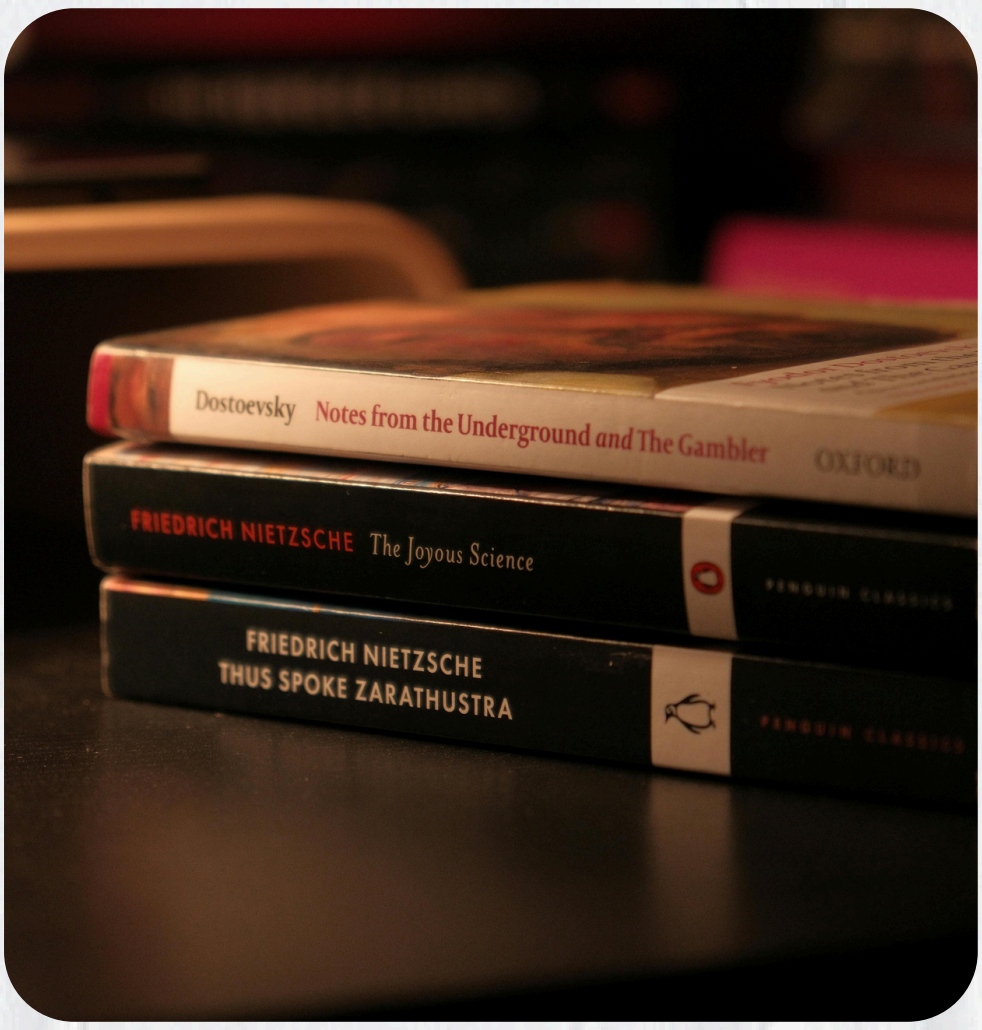

ENGINEERING DISASTERS AND LEGAL ACCOUNTABILITY

On the morning of April 20, 2010, an explosion aboard the Deepwater Horizon drilling rig in the Gulf of Mexico killed eleven workers and triggered the largest marine oil spill in history. In the months that followed, as millions of barrels of oil blackened the Gulf Coast and devastated ecosystems and livelihoods across four states, investigators and lawyers began the complex work of determining what had gone wrong, and who was legally responsible for it. What they found was a catastrophic convergence of engineering failures, safety shortcuts, and organizational dysfunction that would take years of litigation and billions of dollars to adjudicate.

The Deepwater Horizon case is an extreme example of a fundamental challenge at the intersection of law and engineering: when complex technological systems fail with devastating consequences, how does the legal system assign responsibility, establish causation, and deliver justice?

Engineering disasters occur with troubling regularity. The collapse of the Champlain Towers South condominium in Surfside, Florida in 2021, which killed 98 people, raised questions about structural engineering standards, building inspections, and the liability of condominium associations.

The Boeing 737 MAX crashes of 2018 and 2019, which together killed 346 people, exposed fatal flaws in an automated flight control system and triggered sweeping litigation against Boeing and the FAA. The 1984 Bhopal gas leak, each of these tragedies illuminates the profound human consequences when engineering systems fail, and each generated complex legal battles over causation and accountability.

Establishing legal liability in engineering disaster cases requires bridging two very different worlds: the probabilistic, systems-oriented language of engineering and the binary, fault-based framework of tort law. Engineers think in terms of risk distributions, failure modes, design margins, and systems interactions. Lawyers and courts think in terms of duty, breach, causation, and damages. Translating between these frameworks, in a courtroom, under adversarial conditions, in front of a lay jury, is one of the most demanding tasks in all of litigation.

Expert witnesses are the primary vehicle for this translation. In complex engineering cases, both sides typically retain teams of highly credentialed engineers, scientists, and technical specialists who analyze the evidence and offer opinions on questions like: Did the defendant's design meet applicable standards? Did a specific failure cause the harm in question? Was the risk foreseeable? Could it have been prevented at reasonable cost?

The challenge is that engineering judgment is rarely black and white. Standards and codes set minimum requirements, but reasonable engineers can differ about what best practice requires in novel or ambiguous situations. The question of whether a specific failure "caused" a disaster can be technically complex; large-scale failures are typically the result of multiple interacting factors, making simple linear causation difficult to establish. Juries, who may have limited technical background, are asked to evaluate competing expert opinions on extraordinarily complex subjects.

Regulatory frameworks play a critical supporting role. Agencies like the FAA, the Nuclear Regulatory Commission, OSHA, and the EPA set technical standards for high-risk industries and investigate major accidents. Their findings, contained in accident investigation reports, enforcement actions, and regulatory proceedings, often form the backbone of subsequent civil and criminal litigation.

Criminal liability in engineering disaster cases is rarer than civil liability, but not unknown. Following the 737 MAX crashes, Boeing entered into a deferred prosecution agreement with the Department of Justice, paying $2.5 billion to resolve charges of conspiracy to defraud the FAA. Individual engineers and managers have faced criminal charges in cases ranging from the Bhopal disaster to the Flint water crisis, where officials were charged with involuntary manslaughter for decisions that contributed to lead contamination of the city's water supply.

The legal consequences of engineering failures extend beyond compensation and punishment. They shape engineering practice, regulatory standards, and corporate culture. The threat of liability creates powerful incentives for safety, but also, critics argue, can create perverse incentives for defensive documentation and risk aversion that impede innovation.

The relationship between law and engineering is ultimately one of mutual necessity. Law depends on engineering to understand what happened and why. Engineering depends on law to enforce the accountability that makes safety a rational investment. When that relationship works well, the result is a safer, more reliable built world. When it fails, the consequences can be catastrophic and irreversible.