Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Frontier Technologies, Complex Externalities, and the Evolution of Private Order: Implications for AI Governance

Ali Palida

apalida@ihs.gmu.edu

Institute for Humane Studies

Nicholas Thielman

R Street Institute

Research Article

Keywords:

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11369312/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Frontier Technologies, Complex Externalities, and the Evolution of Private Order: Implications for AI Governance

Ali Palida[1]

Nicholas Thielman[2]

Abstract[3]

Artificial intelligence (AI) creates externalities whose sources and magnitude are difficult to know before deployment. Most of the work on AI governance asks how public authorities should regulate these risks. We ask how public regulation affects the discovery of private order. The cases of electrical safety certification, cyber insurance, and early radio broadcast illustrate how experience can reveal conflicts and support flexible private governance. A model of reversible deployment identifies two costs of an imposed public veto. By blocking deployment, it can prevent the learning needed to identify effective safeguards and allocate authority. Because affected parties cannot bargain around it, the veto can also continue to block mutually beneficial activity after they acquire more information. The cases and the model suggest that the costs of public regulation include its effects on private governance that might otherwise emerge. This consideration matters for the design of institutions that govern AI.

1 Introduction

Artificial intelligence (AI) is a general-purpose technology whose applications can create benefits and harms across many individuals. These effects are often difficult to identify before deployment because they depend on the particulars of the environment. This creates a problem of complex externalities: policymakers and private actors alike may not know in advance

[1]Institute for Humane Studies, Arlington, VA.

[2]R Street Institute, Washington D.C.

[3]No funding was received for this research.

which safeguards will prove effective (Paniagua, Rayamajhee, and Murtazashvili 2024). Most work on AI governance asks how public authorities should regulate these risks. We ask a different question: how does public regulation affect the discovery of private governance?

Our argument is that experimentation with a frontier technology can reveal not only its benefits and harms, but also the institutions appropriate for governing them. We develop this argument in two steps. First, we examine three cases—Underwriters Laboratories, cyber insurance, and early radio broadcasting—in which experience with new technologies generated information that private actors translated into various forms of control. The cases illustrate a common process: use reveals conflicts and risks, and governance becomes more specific as knowledge accumulates.

Second, we formalize one mechanism within this process in a model of experimentation and veto rights. In the model, use of a frontier technology generates a multidimensional vector of uncertain payoffs and information about their underlying distribution. Affected parties periodically bargain over private veto rights, which can be reassigned as they learn. A public veto is imposed externally and cannot be removed through private agreement. The technology itself is reversible, so persistence comes from the governance arrangement rather than from technological irreversibility.

The model identifies two costs of public ordering under uncertainty. First, a public veto can suppress learning by blocking the activity that generated information. Second, it can prevent affected parties from acting on information once it is acquired. Private parties can revise their own rights at the next bargaining date; they cannot bargain around a continuing public veto.

Open access appears as one possible arrangement within private ordering, not as a commitment to permanent rulelessness. When private rights can be reconsidered before every opportunity for use, open access has no distinctive advantage. Alternatively, when renegotiation is infrequent, an earlier assignment of veto rights may temporarily block valuable experimentation. Allowing open access can preserve those opportunities, but it can also permit deployments whose private benefits are smaller than the losses imposed on others.

The broader contribution is to connect technological learning to the evolution of private order. Research on self-governance and private ordering shows how governance institutions

can emerge in response to particular conflicts (Ostrom 1990; Leeson 2007, 2009; Skarbek 2011, 2014; Shortland 2017). We emphasize that, at a technological frontier, the relevant conflicts and appropriate institutions may themselves have to be discovered. Governance is not simply a mechanism for controlling known externalities; it can be an outcome of the process through which externalities become understood.

This perspective matters for AI governance. Public intervention may prevent harmful deployments, particularly where expected losses are severe or private arrangements cannot represent affected interests. But regulation also has dynamic institutional costs. By blocking experimentation, it may suppress information needed to develop better safeguards. By fixing authority outside the bargaining process, it may also preempt governance arrangements that private actors would otherwise adopt as knowledge accumulates.

The remainder of the paper proceeds as follows. The next section explains why AI application deployment generates complex externalities. The following section examines private ordering in electrical safety certification, cyber insurance, and spectrum management. We then present the model of experimentation and veto rights. The discussion connects the model back to the cases and considers its implications and limits for AI governance. The final section concludes.

2 Private Governance, Complex Externalities, and AI Safety

The deployment of AI applications can affect individuals and organizations beyond those directly involved in developing or using them. An application may expose third parties to spillovers that propagate through interconnected systems. These consequences constitute externalities insofar as their costs or benefits are not fully reflected in the decisions of those deploying the application.

AI application deployments can generate *complex* externalities in the sense identified by Paniagua, Rayamajhee, and Murtazashvili (2024). First, their effects can reach numerous parties across organizations and jurisdictions, including individuals outside the deploying

firm's contractual relationships. Second, causal sources can be difficult to identify because outcomes depend jointly on environmental characteristics. Third, uncertainty about the magnitude and distribution of harms complicates general regulatory rules: the same model can support applications with different risks, and some consequences become apparent only through use. These characteristics impede both private bargaining and public efforts to specify appropriate controls in advance.

Complexity makes acquiring knowledge central for institutional comparison. Private parties may struggle to negotiate safeguards. Public authorities require much of the same information to determine appropriate policy and regulation. The question is how alternative governance arrangements obtain this information, translate it into effective controls, and revise those controls as knowledge changes.

Research on self-governance establishes that difficult externality problems need not be addressed exclusively through centralized authority. Ostrom (1990) documents how resource users develop rules using knowledge of local conditions. The broader polycentric perspective emphasizes how multiple, overlapping centers of authority can experiment with different arrangements and adjust them in response to experience (Ostrom, Tiebout, and Warren 1961; Ostrom 2010).

The private ordering literature extends this perspective to settings where conventional public enforcement is weak or unavailable. Leeson (2007, 2009) shows how reputation, contractual practices, private adjudication, and sanctions can support cooperation under such conditions. His comparative approach evaluates imperfect private arrangements against feasible governmental alternatives. Skarbek (2011, 2014) similarly demonstrates that governance organizations can emerge in response to unmet demands for public services. Their organizational forms vary with environmental characteristics and the costs of providing governance. Together, these studies suggest that institutional arrangements can develop in response to the particular conflicts and coordination problems that participants face.

Insurance provides another mechanism for translating dispersed information into governance. Shortland (2017) examines how insurers and associated intermediaries developed arrangements for managing kidnapping risks across fragmented jurisdictions and under substantial uncertainty. Organizations exposed to losses have incentives to acquire information,

evaluate precautions, and make coverage conditional on conduct. Insurance can consequently combine risk transfer with monitoring and contractual control. Alongside these benefits, it supplies a potential means of governing technological risks whose characteristics become clearer through experience.

For AI applications, deployment can generate information needed to develop these arrangements. Experience reveals how applications perform, where losses arise, and which precautions work, creating opportunities for parties to devise more specific controls. As Wildavsky (1988) emphasizes, efforts to prevent technological hazards must account for the learning that experimentation makes possible. Governance choices can affect the knowledge available for improving subsequent governance.

The following cases examine this process in electrical safety certification, cyber insurance, and radio spectrum governance. Underwriters Laboratories illustrates the development of testing and standards in response to unfamiliar hazards. Cyber insurance illustrates continuing adjustment of contractual controls as risks change. Early spectrum governance illustrates the emergence of private arrangements and their subsequent political displacement. These cases provide institutional context for the model's narrower analysis of learning about deployment payoffs and revising veto rights. They motivate an approach to AI governance that considers how restrictions affect both external harms and the discovery of institutions capable of addressing them.

3 Private Ordering in Technological Governance

3.1 Underwriters Laboratories and Electrical Safety

While the physical properties of electricity had been discovered and known for centuries, it was not until 1879, when Thomas Edison invented the incandescent light bulb, that these physical characteristics became channelled into consumer products (Energy Information Administration 2007). In 1882, Edison opened the first commercial power station in New York City, which could deliver power to households via direct current, and in 1884 Nikola Tesla invented the electric alternator, allowing for the deployment of alternating current in convey-

ing electricity. It was in 1893 that the Westinghouse Electric Company used an alternating current system to light the Chicago World's Fair, which provided the broader public with its first exposure to the capabilities and general applications of electricity (Ibid.; Brearly [1923] 1997, 75). Electricity is a general-purpose technology (GPT) in that it represents a "core" technology upon which a "cluster of innovations" emerges as both the core and derivative technologies are improved and developed (Schumpeter [1934] 2021; Bresnahan and Trajtenberg 1995). By causing discontinuous jumps in innovation, GPTs are thought to drive both productivity and economic growth.

Nevertheless, the introduction of electricity and electrical appliances posed many challenges, not least because of the hazards they posed to user safety. Initially, fires caused by electrical equipment malfunctions had become an increasingly large loss driver for insurance companies. However, there was a great deal of uncertainty regarding the source of these hazards, that is, which applications of electricity posed greater or lesser hazards to user safety (Lytton 2014, 543). Indeed, the novelty of electrical technologies, combined with their growing pervasiveness, meant that the externalities involved transcended both jurisdictional and sectoral boundaries. As one contemporary historian put it, "one may visualize the entire burnable property of the United States... and try to realize that all of this unimaginable wealth is, in a sense, tied together by a network of insurance contracts that provide security to the holders of more than thirty million separate policies... insurance companies and the millions of policy holders have an identical interest in obtaining accurate knowledge of the elements of hazard that must be considered in fixing rates" (Brearly [1923] 1997, 78). The magnitude of the risks involved created a demand for safety governance, while simultaneously making it difficult for any single jurisdiction or industry to identify and impose uniform standards. Rather than necessitating a public mandate, this uncertainty provided the impetus for entrepreneurs in the insurance industry to devise novel governance mechanisms by which it could be overcome in the form of Underwriters' Laboratories (UL).

Started by MIT-educated electrician William H. Merrill in 1893, UL initially existed as an appendage of the Chicago Underwriters' Association before being contracted by Western Union Insurance to perform its testing and certification services on electrical equipment and building materials for insureds across its territory (Ibid., 76; Lytton 2014, 543). By 1895,

UL's reputation had grown enough that the National Board of Fire Underwriters agreed to fund its operations, and, with this increase in resources, UL began publishing a list of approved devices which was distributed to fire underwriters and municipal fire service officers (Lytton 2014, 544). By 1901, UL was spun off as an independent corporation and had become financially independent by 1916. By then, UL's services had been extended to not only safety testing but also factory inspection and product labelling (Ibid.).

The expertise of UL's staff and the quality of its certification services contributed to the growing demand for its services. As Brearly put it, UL served to adjudicate between the differing opinions of specialists by serving as a "court of last resort such as could be found only in an adequate laboratory, where tests could be made without previous bias and their results could be certified" (77). Insurance providers initially sought UL's services as a means of reducing the financial risks associated with insuring structures containing electrical appliances and implements, the quality of which could not be easily ascertained. These services also came to be demanded by construction companies, municipalities, and consumers themselves, all of which faced similar information asymmetries. Eventually, insurers came to condition underwriting on the use of UL-approved building materials, local governments incorporated their standards into building codes, and manufacturers sought the UL label for their products, owing to its sales value (Ibid., 81; Lytton 2014, 544-545).

Competition from other safety certification organizations, city governments and urban safety councils served to discipline UL's own internal governance, thereby preserving the quality of its brand. Moreover, UL was embedded in a network of other fire safety organizations, generating further reputational incentives to both preserve quality and to share best practices with others in the industry (Lytton 2014, 547). To this end, UL has maintained high internal standards for both personnel and the procedures used for inspection and testing. Under Merrill, UL hired engineers with technical degrees and organized them into separate, specialized departments. These engineers were "practical men who have seen many devices fail under test and know what to look for..." (Brearly [1923] 1997, 81). In addition to this, Merrill held employees to high ethical standards which were overseen, in turn, through a variety of internal controls designed to ensure oversight of inspectors and to minimize the chances of mistakes or misconduct. Put differently, Merrill built redundancy

into the structure of UL, thereby improving the robustness of its certification and inspection procedures by reducing the likelihood that any single point of failure could undermine their quality.

UL's internal governance was further aided by the broader expert ecosystem it was embedded within. Merrill established independent oversight councils to assess the findings of its engineers, often composed of independent insurance experts and government engineers. These bodies served as additional verification points, certifying the test results passed along from internal UL engineers. Complete reports were then shared with clients, and the certified product was then added to UL's published list of approved products. Both of these procedures allowed for further external verification. UL also routinely submitted its standards for recognition as national consensus standards by the American National Standards Institute, which involved a process of public input (Lytton 2014, 545-546). Lastly, Merrill sought to inculcate esprit de corps within UL by emphasizing the "public good" element of their work.

It was these internal and external sources of discipline and monitoring that allowed UL to become the leading safety certification and testing organization in the United States, allowing it to successfully navigate the tumult of the Great Depression and Second World War. The post-war period saw an increase in safety consciousness within the public and in government. Such groups as Consumer Research, Inc. and the Consumer Union made increasing use of UL certifications to help inform the public about the relative reliability and hazards of consumer products in addition to their own investigations and testing (Knowles 2017, 52-53). Yet, it was also during this period that UL found its position in the marketplace increasingly challenged by the federal government. Nevertheless, despite the creation of the Consumer Product Safety Commission and challenges to its non-profit, tax-exempt status, UL adapted and was able to expand its operations. Chief among these was an increased focus on consumer advocacy and engagement with a wider range of stakeholders. Under then President Baron Whitaker, UL added consumer groups, government agencies, public utilities, and safety experts, in addition to the insurance industry, to the groups represented on its board of trustees. This expansion was accompanied by a rule that no single stakeholder could make up more than one third of the board's membership (Ibid., 60). Additionally, UL organized a consumer advisory council with its own dedicated testing lab to better serve the

growing demand for consumer product testing (Ibid., 60-61).

Congress subjected UL to scrutiny in May 1967 over the recall of UL-certified General Electric television sets, which had been discovered to be exposing consumers to radiation. In response, Whitaker oversaw a reevaluation of UL's internal procedures. The changes included increased oversight and rigor in testing and product sampling, as well as follow-up examinations of manufacturers, with the number of inspections required to maintain certification increased from one to four per year (Ibid., 60). Additionally, UL began requiring manufacturers to include its label on any product certified and introduced a "hold harmless clause" which required manufacturers to compensate UL for lost business owing to the inclusion of their label on products that had not met their standards (Ibid., 61). In effect, this clause functioned as a Coasian compensation mechanism whereby manufacturers internalized the costs imposed on UL by misappropriating their label. All of these changes were made to adapt to the changing times and market demands. Owing to these changes, UL not only survived but prospered, doubling its total business over the course of the 1960s and early 1970s. With increased globalization starting in the 1980s, UL saw demand for its services, as well as growing competition from other private standard-setting bodies, helping to construct the rules of the game by which globalization took place. To this day, UL remains a leader in private standard-setting, "dwarfing federal and state-level efforts at safety testing of consumer products and building materials" (Knowles 2017, 62).

Overall, the case of Underwriters Laboratories demonstrates the ability of private institutional entrepreneurs to respond to market failures associated with novel technologies by constructing governance arrangements capable of internalizing complex externalities in an evolving technological environment. It is worth emphasizing that UL's internal governance and quality were not driven purely by the public-spiritedness of its founder or subsequent leadership. Rather, these measures reflected the competitive marketplace that UL existed within and, particularly after 1916 when it became reliant upon charging fees for its service, the need to earn positive profits to survive (Alchian 1950). If it failed to "certify the attainment of definite and adequate standards" in its product inspection or safety testing (Brearly [1923] 1997, 80), the consequence would be a loss of business as manufacturers, insurers, and municipalities turned to other safety certification providers. Thus, investments

in signal credibility were as much a competitive necessity as a product of the entrepreneurial vision of its leadership. Moreover, this competitive environment was defined by "heterogeneous interacting elements produc[ing] the same or similar functions in the system through a variety of methods." That is, it embodies what Roger Koppl (2018, 184-185) refers to as "synecological redundancy." UL exists in an international marketplace where it competes not only with other private certification providers but also with government agencies, insurance companies, and nonprofits that provide overlapping safety-governance functions. It is left to the market and consumer preference to adjudicate which of these institutional forms provides the most trustworthy signal of safety and reliability.

In the context of AI governance, various scholars have advanced proposals that adapt elements of the "UL model" in the form of Independent Verification Organizations (IVO) operating in regulatory markets (Hadfield and Clark 2026; Ball 2025). These proposals involve the federal government establishing some set of standards or outcome-based safety requirements which would then be translated into practical governance by competing government-licensed auditing bodies. The case of UL, however, suggests that the discovery process need not begin with the government establishing the substantive safety objectives against which private organizations compete to provide verification. Indeed, the case of UL illustrates that private institutional entrepreneurship can operate in the discovery and selection of safety standards and requirements as well. The imposition of a regulatory market framework on AI governance would effectively constrain the experimentation and selection of these standards to objectives chosen by the government, potentially forgoing more effective or efficient standards.

Historical experience and insights from basic public choice theory suggest that democratic legislatures are unlikely to reliably identify what appropriate standards or metrics licensed governance organizations should aim at, nor are they likely to update these in an efficient manner (Ben-Shahar and Logue 2012, 198-199). The legislating of safety metrics would likely invite industry influence. Indeed, the issue of AI safety seems particularly vulnerable to "Bootleggers and Baptist"-style regulatory capture (Metcalf 2026; Thielman 2026). While proponents of the IVO model are aware of this risk with regard to government agencies determining safety rules, their approach leaves a related political-economy problem at

the legislative stage, where special interests, or the interests of politicians themselves, may influence the substantive objectives that private regulators are subsequently tasked with implementing (See, for example, Ball 2025, 11-15). There exists also the issue that legislatively determined rules could privilege one particular perspective on AI safety by, for example, ignoring cybersecurity experts or perspectives informed by experience with other high-risk technologies like nuclear fusion.

Developing adequate safety standards for AI requires that the experimentation and synecological redundancy that produced modern safety standards and testing be preserved by not constraining it around some legislatively determined baseline. These considerations suggest that, insofar as public authorities play an active role in determining these standards, it should be limited to the establishment of voluntary standards or guidance (Thierer 2023). In this way, public agencies function as another participant in the AI governance marketplace, rather than setting the rules by which competition takes place, allowing the market process to develop and adapt standards as conditions change. Otherwise, private actors and competition are capable of discovering and setting credible safety standards.

3.2 Cyber Insurance and Private Governance of Cyber Risks

Cybersecurity has traditionally been considered a public good by policymakers and thus not amenable to private production and exchange. Investments in cybersecurity are said to embody both positive and negative externalities (Hodgins 2025a). One firm's decision to invest in security entails positive externalities which are not fully internalized, creating incentives for other firms to free-ride off the investments of more secure firms. As a result, firms acting independently will tend to underinvest in cybersecurity systematically. This poses additional issues, as individual underinvestment can leave security gaps that, given the interconnected nature of modern digital systems, create vulnerabilities for other firms and individuals. That is, underinvestment creates negative externalities for other users (Moore 2010, 107). The potential for underinvestment by actors spread across digital ecosystems makes the actual distribution of risks difficult to quantify. Moreover, the interconnectedness of these externalities means that the number of impacted parties is potentially substantial. The Federal Bureau of Investigation, for example, estimated that the Salt Typhoon hacks affected as

many as 600 companies across 80 countries (Viswanatha and Krouse 2025). Thus, the scope of cyber externalities can be large, impacting parties spread across numerous jurisdictions, thereby increasing the coordination costs of creating uniform governance architectures.

Lastly, cyber threats represent dynamic risks. The threat surface changes with the technical capacities of both attackers and defenders, meaning that the relevant risks are not merely difficult to measure but can themselves change as actors experiment with new technologies and techniques. The July 2026 Hugging Face incident was subsequently discovered to have been perpetrated by a coordinated "swarm" of 700 AI agents that used an unsanctioned message board to break out of an OpenAI test environment (Greenblatt, Cotra, and Wijk 2026). These dynamics are compounded by information sharing problems: firms that suffer an attack may face incentives not to share this information owing to the potential market or reputational penalties that may be incurred (Powell 2005, 502; Eling, Elvedi, and Falco 2023, 429). Owing to the interconnected nature of cyber-related externalities, combined with the difficulty of assessing ex ante where threats will manifest in digital ecosystems and their evolving nature, cybersecurity represents a complex externality.

The collective nature of cybersecurity has been offered as justification for public interventions designed to overcome market failures in the provision of cybersecurity. Nevertheless, scholars such as Powell (2005), Coyne and Leeson (2005), Dourado (2012) and Hodgins (2025a, 2025b) have questioned the efficacy of such an approach. The uncertainties that plague the private provision of cybersecurity are shared by public provision as well. Simply put, government agencies lack the means to effectively calculate what the optimal amount of cybersecurity is and the feedback mechanisms necessary to adjust this quantity as threats evolve (Powell 2005, 505). Moreover, investment in cybersecurity measures entails significant private benefits directly to the user; thus, firms and individuals do have an incentive to make these investments independently of others (Powell 2005, 500; Dourado 2012). Indeed, the growth in cyber insurance represents one organizational form by which private ordering can align incentives for cybersecurity investment by linking firms' security practices to the cost and availability of insurance. As such, cyber insurance has allowed for more fluid governance, allowing the rules of the game to evolve with the threat environment.

Similar to the evolving AI governance framework, cyber governance exists in a poly-

centric institutional context. A number of sectors, such as healthcare, financial services, and energy, are subject to regulations regarding cybersecurity. Under the Health Insurance Portability and Accountability Act of 1996, for example, the healthcare industry is required to implement certain security measures to protect consumers' personal health information. Financial institutions are subject to similar requirements to safeguard the personal data of their customers under the Gramm-Leach-Bliley Act. While these regulations require firms to take certain positive steps to protect user data, or ensure minimum security measures are implemented, they function more as basic ground rules than prescriptive regulations. Moreover, as will be discussed below, public cybersecurity regulation suffers from several pathologies less pronounced in private cyber governance.

Cyber insurance emerged in response to the demand on the part of firms for a means to offset the financial liability associated with cybersecurity incidents (Herr 2021, 106). While certain kinds of cyber risks, such as computer fraud, had been covered by Lloyd's and other large insurance providers, it was not until 1997 that the first dedicated cyber insurance policy was introduced by AIG (Woods and Wolff 2025, 3). Early cyber insurers faced a first-mover problem in that there was uncertainty over the kinds of risks to be covered, their scale, as well as how best to balance their exposure to these risks. The coverage provided by conventional cyber insurance has evolved with changes in loss drivers, shifting from a focus on indemnifying policyholders for the business downtime resulting from cyberattacks to a more tailored focus on the risk of data breaches resulting from ransomware and social engineering (Woods and Wolff 2025, 3-6). From the late 1990s through the early 2000s, the market for cyber insurance grew slowly. By 2012, for instance, the U.S. cyber insurance market was only valued at $2 billion (Herr 2021, 99). However, as the cyber threat environment grew, along with the broader societal awareness of cybersecurity, so too did the demand and supply of cyber insurance products (Herr 2021, 106). Between 2017 and 2022, the cyber insurance market grew by approximately 30 percent worldwide, although this growth slowed to 5 percent annually between 2020 and 2025. The dollar value of global written cyber insurance premiums rose from approximately $3.5 billion to $15.6 billion over this same period (Tashima 2026). As the scale of risk transfer has increased, so too has the incentive of insurers to monitor and enforce proper cybersecurity procedures on the part of

policyholders and adjust premiums accordingly. Furthermore, cyber insurance entails a risk of moral hazard in that insureds may increase risky behaviors, thus expanding the potential liability of insurers (Arrow 1963; Hölmstrom 1979).

Ben-Shahar and Logue (2012) note that insurance companies possess several instruments by which to overcome the difficulties posed by moral hazard in addition to the traditional tools of setting risk-rated premiums, deductibles, and copayments, including standard setting and partnering with specialist firms to overcome information asymmetries. Early cyber insurers relied upon a combination of in-person audits and actuarial modelling to assess premiums and to gauge the security postures of policyholders. While crude, such methods created incentives for insureds to attend to their security protocols because such efforts were rewarded by insurers. AIG, for instance, offered a 25 percent discount to clients whose systems were audited and certified by the National Computer Security Association (Woods and Wolff 2025, 4). In performing this auditing function, cyber insurers effectively incorporated a set of security standards into their policies, providing external market participants with a credible signal by which to assess the posture of policyholders (Herr 2021, 103).

However, as digital technologies developed, insurers faced recurrent issues with data availability, continually relying on a combination of surveys, data from public breaches, and knowledge drawn from industry practice and experienced underwriters (Ibid., 4). Moreover, as cyber incidents increased, so too did demand for insurance products, driving industry profitability. The presence of economic profits encouraged entry, as insurers saw opportunities in offering specially tailored cyber insurance products as compared to other product lines; the median loss ratio in 2018 for cyber insurance was 35 percent as compared to 62 percent in property and casualty insurance (Ibid., 4). Woods and Wolff contend that increased market entry limited the ability of insurers to collect information from customers and enforce standards; however, this remains debatable. Insurance market competition generally rewards insurers that "encourage their insureds to adopt good risk management practices," since firms that fail to induce efficient care risk losing business to firms that demonstrate superior risk governance (Ben-Shahar and Logue 2012, 204). Identifying better risk governance standards allows insurers to charge a lower premium and thereby expand market share as well as reduce the risk that claims costs outrun premiums (Ibid.). Herr (2021)

notes that enforcement of security standards does not come exclusively from insurers, but from a polycentric process occurring through the market itself. The signal provided by a firm's adherence to particular standards was valued both by insurers and by the business partners of the insured firm located both up and down its supply chain. Indeed, the demand for cyber insurance products and their associated standards reflected both a requirement of doing business and a means to shift some of the financial risk associated with cyber incidents (Ibid., 106-107). The accountability imposed by market participants collectively provided insurers with both the legitimacy and the incentives to enforce the standards embedded in cyber insurance contracts. Firms that failed to maintain adequate security, obtain insurance, or both risked reputational and financial penalties.

Regardless, by 2018 the locus of cyber threats had shifted toward ransomware and social engineering, further increasing the complexity of cyber protection and underwriting for insurers. In response, providers began partnering with or acquiring InsurTech firms serving as managing general agents, thereby gaining the expertise necessary to assess the riskiness of clients. Other insurers partnered directly with cybersecurity firms and cloud providers, gaining direct access to customer data. The increased demand among small and medium-sized enterprises for cybersecurity products also helped with overcoming data gaps. Together, these trends assisted insurers in modeling and pricing the risk of attacks into premiums (Woods and Wolff 2025, 5-6). Enhanced data availability allowed insurers to tailor standards enforcement to particular clients, with some insurers increasingly requiring that firms adopt multi-factor authentication or partner with companies like Google and Microsoft, which provide such services as a condition for obtaining insurance. In effect, through more efficient monitoring, insurers can "try to guide policyholders towards improved security practices" (Ibid., 6).

Reinsurers and capital market participants are embedded within the cyber insurance ecosystem, functioning to buttress the standards and enforcement imposed by primary insurers. As the market for cyber insurance expanded throughout the 2000s and 2010s, insurance providers turned to reinsurers to backstop their own liabilities. Reinsurers collect information from across the primary market, allowing them to monitor and adjudicate different underwriting policies, rewarding insurers that demonstrate superior risk management.

Alternatively, riskier insurers face higher premiums, lower coverage limits, heightened monitoring, and, in the last resort, rejection of service. In effect, reinsurers serve as both claims repositories and as a source of market discipline on primary insurance providers (Ibid., 8-9).

The role of capital markets in cyber insurance has been to serve as another avenue by which insurers and reinsurers offload financial risk via catastrophe bonds, as well as another layer of monitoring and market discipline. Starting in 2016, Credit Suisse issued the first series of catastrophe bonds that included cyber risks as part of their coverage. It would not be until 2023 that dedicated cyber catastrophe bonds would enter the market. Much like other catastrophe bonds, these instruments promise investors a return above that offered by conventional bonds in exchange for the risk that, should certain pre-defined contingencies occur, the funds invested will be surrendered for use by the issuing insurer (Ibid., 9-10). This risk incentivizes monitoring by bond investors of the underwriting practices of insurance companies. Woods and Wolff offer a plausible mechanism by which capital market discipline encourages prudent risk management by cyber insurers: "Rational markets would offer the capital most willingly to those insurers who can demonstrate superior underwriting processes. Cheap capital allows those insurers to expand underwriting. This could improve social welfare if capital markets reward insurers who actively improve the security posture of policyholders" (11). Indeed, there is evidence that capital markets are increasingly working cybersecurity into their assessments of not only insurers, but of firms generally. As Fitch recently stated in separate reports on healthcare and water utility service providers, those "most likely to mitigate potential negative ratings momentum are those that demonstrate business resiliency and have robust incident response capabilities, effective operational continuity planning and sufficient ratings headroom to absorb the financial stress of a cyber event" (Geller 2026).

Taken together, the layered market discipline exerted by insurers, reinsurers, and capital markets represents an example of how markets and private governance are currently addressing complex externalities through polycentric arrangements that evolve in real time. Early insurers experimented not only with methods of underwriting cyber risks, but with the very scope of the risks that insurance should cover. Despite this initial uncertainty, entrepreneurs developed products for covering different cyber risks and methods for auditing and monitor-

ing insureds. By allowing experimentation in security procedures and creating an incentive-compatible mechanism to enforce and monitor them, cyber insurance enables learning that allows entrepreneurs to adjust the institutions governing cyber risks to an evolving threat environment. This process of experimentation is particularly important because, as Woods and Wolff note, when it comes to reasonable or adequate cybersecurity, "there is no single definition, but instead a range of contextual definitions" (p. 12). The robustness of this private ordering arrangement is further indicated by the finding that under modeled extreme loss scenarios resulting from hacks of critical infrastructure, cyberattacks remain within the realm of private insurability (Eling, Elvedi, and Falco 2023). At technological frontiers characterized by uncertainty and complex externalities, decentralized institutional arrangements can facilitate experimentation, information discovery, and adaptation. Centralized intervention risks crowding out such private ordering institutions, thereby suppressing the discovery process that gives rise to more efficient risk governance.

As noted above, cyber governance exists in a polycentric institutional environment, where public cybersecurity regulations exist in the background of private efforts to monitor, price, and mitigate ongoing cyber risks. While cyber governance remains an evolving policy environment, there is nevertheless reason to believe that cyber insurers may be comparatively more efficient providers of governance than their public counterparts. Powell (2005) notes that public agencies tasked with cybersecurity are not well positioned to calculate the optimal amount of cybersecurity. While private markets may suffer from underprovision of cybersecurity, the incentives faced by bureaucrats likely bias decision-making toward overprovision. Enhanced cybersecurity can serve as justification for expanded budgets; moreover, given the risk aversion of regulators, they have strong incentives to increase cybersecurity requirements to minimize potential reputational fallout following a cyberattack (p. 505).

A recent report from the Government Accountability Office (GAO) provides evidence consistent with this tendency toward bureaucratic overshooting. The GAO notes that, of the 117 cybersecurity regulations it identified as issued by 37 different federal agencies, approximately 70 percent potentially subject firms to duplicative reporting requirements concerning "cybersecurity incidents, cybersecurity plans or other technical information, or reviews, audits, or assessments" (Government Accountability Office 2026, 4-6). At a GAO-

convened roundtable, industry representatives stated that complying with these regulations consumes time and resources that would otherwise have been spent on responding to incidents or improving their security postures (Ibid., 4). Absent the feedback provided by profit-loss signals and competition, government agencies simply have no reliable way of discovering whether the marginal cyber requirement is "too much" or "too little." Indeed, even if the optimal level could be identified by other means, agencies lacking the discipline imposed by market competition face little incentive to correct toward this level.

Cyber insurers, by contrast, face market pressures to ensure that the governance they provide is relatively efficient. Insurers that fail to identify efficient risk reduction mechanisms risk losing business to firms that succeed in this. Profitability, comparisons with peer firms, and claims experience are among the feedback mechanisms by which insurers can determine whether premiums are too high or contractual standards too onerous (Powell 2005, 505). Moreover, through the underwriting process, partnering with Insurtech firms and other managing general agents, insurers have access to information that is unavailable to government agencies. All of this means that cyber insurers have better information and stronger incentives to govern cyber externalities efficiently. In this way, insurance premiums can perform much the same function as "Pigouvian" taxes, pricing in the value of externalities and altering insured behavior where necessary (Ben-Shahar and Logue 2012, 233-234).

Importantly, from the standpoint of operating on or near the technological frontier, cyber insurance preserves optionality for insureds by presenting a menu of safety options at various risk-adjusted prices (Ibid., 231). Firms can compare these and select those options that best fit their own cost structure. In subsequent periods, insurers can assess firms' risk posture and attempt to influence their behavior by adjusting premiums or safety standards, incentivizing insureds to invest in safety improvements where these are economically feasible. Insurers can also "guide" firms toward better cyber hygiene practices and, in the last resort, enforce exclusions or engage in ex post underwriting (Ibid., 210-211, 215-216). By contrast, government safety regulation tends to operate as an up or down decision, leaving firms with fewer opportunities to experiment with different risk levels. As a result, cyber governance by insurers preserves the experimentation that can yield valuable information as well as innovations, while still acting as a brake on risk-taking.

The case of cyber insurance illustrates how insurers, bolstered by capital markets, are able to adapt, price and develop flexible governance arrangements specially attuned to the risks associated with a rapidly evolving technological landscape. Insurance markets can perform an analogous role in AI governance. As Szpruch et al. (2026, 4) argue, by "conditioning coverage on documented assurance artifacts, continuous monitoring, and alignment with emerging standards, underwriters can foster market-led convergence on what constitutes acceptable risk management." In so doing, AI insurers would perform the same function as cyber insurers in helping to develop, adjudicate between, and enforce standards of reasonable security and risk management. Moreover, since insurers are exposed to losses and competitive pressures in performing these functions, they face a strong incentive to "get the rules right," correcting and adjusting them as circumstances change (Weil 2026). Indeed, given the recent Hugging Face hack, cyber insurers are likely already providing a form of AI governance by pricing firms' cybersecurity postures or updating their policies. Coalition, a managing general agent, for example, revised its cyber insurance contract to cover losses resulting from AI system failures and AI-enabled fraudulent wire transfers (Romanosky and Robinson 2026, 18).

Insurance scholars and practitioners have furthermore suggested that many of the catastrophic risks associated with AI, particularly those related to nuclear, biological, and chemical weapons, could be addressed through the creation of an AI mutual insurance company composed of frontier labs. This organization could then farm out assurances to the growing ecosystem of AI evaluation companies such as METR, Accenture, or Vals AI to both ensure lab compliance with good risk management practices and adjust premiums (Trout, Kvist, and Dattani 2026). Yet, despite these developments, insurance markets are still in the early learning period with regard to AI. Two differing approaches to insuring AI have emerged, with some carriers excluding AI related losses from policies and others embracing these risks through the offer of dedicated policies. Starting in the summer of 2025, a number of legacy insurance companies began excluding losses incurred from AI in their general commercial liability lines, particularly from those involving broad coverage and high policy limits (Romanosky and Robinson 2026, 16). Meanwhile, several dedicated AI insurers have emerged such as Testudo, Armilla, and the Artificial Intelligence Underwriting Company,

offering specialty liability policies and, with regard to the latter two, guarantees of model performance and safety (Ibid., 19). Given that the majority of insurers are currently taking a "wait-and-see" approach to AI, the information provided by these initial experiments will likely prove valuable in determining how the insurance industry addresses AI risk (Ibid., 20).

The case of cyber insurance illustrates that, even in the face of a rapidly evolving risk environment, private actors are capable of developing policies and contracts that price, insure, and govern these risks in increasingly sophisticated ways. The alternative to private governance is not an ideal governance arrangement, but rather a form of public governance with its own flaws and associated costs and risks. There are clear advantages to private governance of cyber risk, as the preceding cases illustrate for other emerging technologies. Yet it is not clear that these beneficial arrangements would have emerged had cybersecurity governance been centrally specified from the outset. With regard to AI, the important point is that these evolving arrangements should not be preempted prematurely, not because private governance is perfect, but because the learning process through which more effective AI governance emerges cannot proceed if it is foreclosed before it can mature.

3.3 Spectrum Management before the Radio Act of 1927

Conventional wisdom holds that before the passage of the Radio Act of 1927 and the creation of the Federal Communications Commission, utilization of the radio spectrum had fallen victim to the "tragedy of the commons." Numerous independent radio stations sought to broadcast to end users in an uncoordinated manner, leading to widespread signal interference issues, wave-jumping, and myriad other externalities. It was, as a U.S. Navy report put it, "etheric bedlam" (Quoted in Coase 1959, 2). In fact, the experience of early radio broadcasting proves this conventional wisdom to be in error; the initial development and commercialization of the radio spectrum took place in a polycentric environment in which private ordering mechanisms played a significant role in allocating and managing scarce spectrum. Indeed, the functioning of private ordering does not imply the absence of law; rather, it implies that state formal governance is lacking or of uncertain effectiveness, or that the state does not directly dictate and enforce the law (Bustamante et al. 2020, 866).

The earliest radio broadcasts consisted of simple point-to-point transmissions between

stations, the number of which was sufficiently small to render interference issues a minor issue. Where these emerged, station operators were quite capable of resolving them through bilateral negotiations. The legal environment of this activity received its first major government intervention following the tragic sinking of the RMS Titanic, the severity of which was partially blamed on the fact that the ship's distress signals were interfered with by signals from amateur radio operators. As a result, the Radio Act of 1912 was passed, which mandated that priority be given to ship distress signals and imposed the requirement that radio broadcasters be issued permits by the Secretary of Commerce (Hazlett 2017, 3; Sidak 2016). Nevertheless, the commercialization of the spectrum did not get underway until 1920 when KDKA began broadcasting from a Westinghouse-owned station in Pittsburgh. From this point forward, the number of spectrum users increased substantially, rising to 600 active radio broadcasts by 1924. This rapid development posed a dilemma, however. A signal broadcast from one station could, if it came within a certain number of kilocycles, confound those from other stations, preventing end users from enjoying clear reception of either due to the interference. This externality problem was further compounded by the growing number of users of the spectrum, increasing the transaction costs of establishing broad-based interference protocols.

Contrary to the conventional story that the complexity entailed in this externality issue could only be resolved via federal regulation, a highly decentralized system evolved for both allocating and managing interference in the spectrum. Thomas Hazlett's work provides an alternative history of the pre-1927 period, showing that "there existed a very lively market in broadcast properties, sold with frequency rights attached, early in the development of the industry..." (Hazlett 1990, 143-144). Hazlett musters an array of evidence to show that before the passage of the Radio Act of 1927, the use and allocation of the radio spectrum was governed in an orderly manner by a system based on common law principles and property rights. Nevertheless, his account also provides evidence that private ordering played a significant role in spectrum governance as well. This fledgling spectrum market was based upon a system of private property rights that had emerged under the formal aegis of the Department of Commerce. As noted above, the 1912 Act empowered the Commerce Department to issue broadcasting permits to minimize interference between users. In practice, this took the form

of recognizing the de facto property rights of existing users through the application of the common law principle of "first come, first served" (Ibid., 147). The assignment of property rights under such a claim rule can, under certain circumstances, incentivize competitive races on the part of claimants that dissipate the value of the unassigned resource (Lueck 1995). However, with spectrum rights, this rule served as the basis for an efficient system of property rights because the then Secretary of Commerce, Herbert Hoover, restricted the assignment of permits, forestalling the competitive dynamics that would lead to dissipation. In effect, this early spectrum market was premised on privately organized claims to particular wavelengths which were then recognized by the Department through the grant of a permit.

Nevertheless, the legal foundations of this system did not represent a "political equilibrium" and were subject to constant challenges (Hazlett 1990, 146-147; Hazlett 2020, 21). The parties opposed to this arrangement consisted of successful commercial stations and politicians, including Hoover. Respectively, these two parties sought to restrict entry into the fledgling broadcasting marketplace and to assert political discretion since "spectrum rents were being capitalized by private owners and, hence, being sacrificed by Congress" (Hazlett 1990, 142, 152-163). Throughout the early 1920s, a succession of radio bills were introduced but failed to gain traction, while suits were filed challenging the Department's ability to enforce frequency assignments. Thus, the state's legal authority to enforce spectrum property rights, and even the nature of these rights, were uncertain. Given this uncertainty, private ordering provided a stable basis by which spectrum markets could develop. Indeed, Commerce officials noted the role of private adjudication in interference disputes between stations. Hazlett quotes a Congressional hearing in which the Department's Solicitor General, Judge Stephen Davis, described how "stations which are operating on one wavelength have been able to get together and agree among themselves" on how best to divide the broadcast space (Hazlett 1990, 145). While Commerce's enforcement powers nominally stood behind these Coasian bargains, Davis noted that the need had not arisen for the Department to intervene to address a dispute (Ibid.).

Elsewhere, legal scholar Frank Rowley similarly noted in a 1927 article that, "there was some interference, due to the surplus of stations over the number of available channels, but in almost every case, station owners showed a willingness to cooperate in making beneficial

adjustments" (Rowley 1927, 5). Beyond bargaining, another avenue by which property rights were reallocated was through the buying and selling of broadcasting stations, and their associated broadcast rights. Given the ambiguity of the Commerce Department's power to police the spectrum, it seems likely that such private contracting served as the institutional centerpiece of this system.

Bilateral contracting, in providing a means by which spectrum users could address externality issues, provided a stable basis on which the fledgling radio industry could develop. Hazlett presents data from the Statistical Abstract of the United States on radio set sales during the first eight years of radio broadcasting, showing that the number of households with a radio set reached approximately 2,000,000 by 1925, before rising to 4,000,000 households in 1926. This increase is quite impressive, particularly given the fact that a new set would cost $1180.72 in 2026 dollars (Hazlett 2020, 27-28). Meanwhile, that same year, the Radio Corporation of America (RCA) introduced the first nationwide radio broadcast via its newly created subsidiary, the National Broadcasting Company (NBC), soon to be followed by a rival network owned by the Columbia Broadcasting System (CBS). These investments produced substantial value to consumers, as indicated by the prices paid by broadcasters for permits in secondary markets. For example, in 1926, WEAF, a New York City-based station, was bought from AT&T by RCA for approximately $1 million, of which $800,000 was paid for its associated spectrum rights (Ibid., 21, 28). Indeed, if the interference issues entailed by the number of users on the spectrum had proven insurmountable, such investments on the part of households and businesses would not have taken place.

This system of private ordering proved robust during a period of quasi-anarchy that followed the 1926 decision in United States v. Zenith Radio Corp. Prior to this, the 1923 Intercity decision had determined that the Secretary of Commerce was not empowered by Congress to ration broadcasting permits, but did have power over times of operation and frequency. Despite this decision, Hoover had continued the practice of issuing permits to users that had demonstrated beneficial use of the airwaves. However, the 1926 Zenith decision overturned Commerce's remaining legal authority over spectrum, which Hoover opted not to challenge (Ibid., 22). What followed was a period where the government's legal authority to police the spectrum was essentially abrogated, allowing some 200 new radio stations to

enter the airwaves, creating myriad interference issues (Ibid.). A novel test case emerged in the case of Tribune Co. v. Oak Leaves Broadcasting Station, Inc. (1926), decided in an Illinois state court, which did determine that "priority in time creates a superiority in right" in spectrum (216, 219). Hazlett argues that from this could have emerged a set of common law rules with regard to property rights. Nevertheless, such a precedent was not yet the law of the land. We offer a slightly different interpretation of this period from Hazlett in that, as the legal regime remained in limbo, private ordering filled the governance void. Indeed, evidence presented by Hazlett lends support to this interpretation.

During this brief period, incumbent spectrum users were reliant on private ordering to enforce property rights and settle disputes with "wave jumpers." Indeed, stock return data for RCA, the largest purveyor of radio sets at the time, indicates that this period of anarchy was more orderly than the conventional narrative implies, with the company's valuation remaining around $50 million throughout this period (Hazlett 2020, 29). This seems to suggest that, even with an increase in the number of users on the spectrum and heightened interference issues, private adjudication provided a reasonably stable basis upon which the wireless industry could continue to function. Had it been otherwise, this would have been reflected in RCA's stock price. Indeed, Hazlett finds that news of developments in spectrum policy, both ostensibly negative and positive, fails to have a significant effect on its stock price, concluding that "this is at odds with the claims for 'etheric bedlam' ... that could only be fixed by centralized authority in the use of radio spectrum rights in wireless services" (Ibid., 32). The passage of the Radio Act of 1927 substituted a system of centralized control, ostensibly to address the complex interference issues entailed by spectrum usage. As Hazlett's work demonstrates, however, rather than addressing a legitimate failure of private governance, the 1927 Act preempted an emerging common law property rights regime that was developing credibly, at the behest of incumbent broadcasters seeking to foreclose competition and politicians seeking to capture spectrum rents.

Nevertheless, the experience of spectrum management before this Act demonstrates that private ordering mechanisms premised on establishing property rights and Coasian bargaining can effectively manage complex externalities associated with novel technologies. As the discussion above illustrates, early spectrum users had significant amounts of their own prop-

erty and other assets at stake, which provided them with a strong incentive to "get the rules right." This decentralized governance process functioned reasonably well as a basis for the development of wireless technologies.

That this system was eventually displaced was not due to its inefficiency or instability. Rather, it resulted from a desire among incumbent broadcasters and political elites to appropriate scarce spectrum rents rather than leave them in private hands. The system the 1927 Act constructed imposed a command-and-control system for allocating spectrum rights centered in the Federal Communications Commission (FCC), premised on the necessity of managing interference issues in the "public interest" (Hazlett 2017; Brough 2025). Far from solving these interference issues efficiently, this centralized regime "replaced a market test with administrative mandate, short-circuiting key feedback loops," introducing a host of new knowledge and incentive problems that hindered the productive utilization of spectrum (Hazlett 2017, 257). FCC bureaucrats ran into the issue faced by all attempts to centrally allocate economic resources: the problem of calculating the value of alternative uses of spectrum. Simply put, the FCC lacked the knowledge to determine ex ante which application of scarce spectrum was best from the standpoint of consumer value. Lacking a profit-loss signal, the FCC relied upon a variety of administrative allocation techniques to determine allocations, including comparative hearings and lotteries. Each carried its own administrative inefficiencies and was marred by rent-seeking by competing spectrum users (Ibid.).

The introduction of spectrum auctions failed to fully address these issues. While auctioning off exclusive property rights to particular bands allows for markets to allocate spectrum to its highest uses, the FCC lacks residual claimancy over the assignment of these rights. Thus, it faces weak incentives to ensure that the structure of these rights is efficient: "In their efforts to avoid border disputes, [the FCC] often separate users into spacious barren silos... this avoids one source of "cacophony," but it overcorrects by imposing a more potent and costlier silence" (Ibid., 313; Bustamante et al. 2020, 864). Moreover, in auctioning off these rights, the FCC is incentivized to maximize revenues obtained at auction or pursue policy goals unrelated to the productive mobilization of spectrum. Such policies, however, allow the FCC to please its political overseers and earn praise from the press. Simply put, "government lacks proper incentives to manage spectrum efficiently" (Beard et al. 2014, 290).

Owing to these issues, the FCC's management has been slow to adapt to changes in wireless technology and the opening of new uses of spectrum. Cellular technology, first developed in 1946, was hindered by FCC policy (Hazlett 2017, chapter 13; Brough 2025). Although difficult to quantify with precision, the foregone benefits from more efficient resource allocation, together with the costs of delayed wireless innovation, have been substantial—incurred, as Hazlett notes, "at the cost of barely perceptible conflicts" (Hazlett 2017, 313).

The shift from decentralized experimentation and private governance in spectrum occurred not for reasons of efficiency, but because this institutional structure failed to achieve a political equilibrium. As a result, it was supplanted by a powerful coalition that erected a centralized allocation regime aimed at limiting entry and rent extraction to the benefit of incumbents and politicians. What benefits this centralized regime achieved in reducing interference among disparate spectrum users were more than outweighed by the costs to innovation, market efficiency, and economic growth. Ronald Coase noted that the purpose of spectrum governance should not be to minimize interference but rather, "The aim should be to maximize output" (Coase 1959, 27). Pre-1927 spectrum markets illustrate that through emergent property rights and bilateral contracting, decentralized spectrum users were capable of mitigating interference issues without public intervention or mandates. Indeed, private contracting has continued to function as an effective means for addressing complex interference problems in modern, liberalized cellular markets (Hazlett, Palida, and Weiss 2024). Moreover, research has shown that self-governance remains a viable mechanism whereby stable property rights can be established that mitigate interference issues (Bustamante et al. 2020).

Turning to AI, the risk is not that private order institutions will fail to emerge to govern AI safety, but rather that premature public intervention will foreclose such mechanisms before these can fully mature. Premature intervention forecloses not only potentially valuable AI applications but also the experimentation through which optimal governance arrangements are discovered. Coasian contracting represents only a single version of such arrangements. As the previous case studies illustrated, private safety certification, insurance contracts, and standards are other private governance mechanisms by which the risks associated with AI can be controlled or mitigated. Indeed, an emergent private governance regime for AI would

likely entail a mix of various mechanisms. What mixture of private governance mechanisms, along with what, if any, role there is for expanded public authority, is optimal remains unknown. The essential point is that the discovery process by which this mix is discerned should not be foreclosed.

4 A Model of Experimentation and Veto Rights

The historical cases show governance developing as people learn what a technology does. Use provides evidence about which precautions work. We model a frontier technology whose effects are initially uncertain in order to isolate how veto rights shape this process. The main comparison is between private rights that affected parties can periodically revise and a public veto that they cannot bargain away. Use is reversible, so persistence in the model comes from the governance arrangement rather than the technology.

4.1 Environment and learning

Time is infinite, with periods $t = 0, 1, 2, \ldots$. There are K risk-neutral parties affected by a potential use of the technology. Each party has the opportunity to initiate use in each period. If use occurs, party i receives payoff x_{it}; otherwise, every party receives zero. The realized payoff vector $x_t = (x_{1t}, \ldots, x_{Kt})$ is fully observed. An application undertaken in one period can be stopped in the next. The use is reversible.

Let θ be the unknown vector of expected payoffs. Conditional on θ, payoff vectors are drawn independently over time according to $x_t \sim \mathcal{N}(\theta, \Sigma)$. The parties share a normal prior $\theta \sim \mathcal{N}(m_0, C_0)$ and discount future payoffs by $\delta \in (0, 1)$.

Use generates both a current payoff and information. At belief state $s_t = (m_t, C_t)$, observing x_t yields

$$C_{t+1} = (C_t^{-1} + \Sigma^{-1})^{-1}, \qquad m_{t+1} = C_{t+1}(C_t^{-1} m_t + \Sigma^{-1} x_t). \tag{1}$$

If use does not occur, no payoff is observed and beliefs remain unchanged. Information about an application arrives through experience with it.

4.2 Private and public veto rights

At dates $0, R, 2R, \ldots$, the affected parties bargain over a private veto set S. Members of S can withhold consent from use. The parties can reassign these rights and make transfers at the next bargaining date, but S remains fixed for the intervening R periods. The interval represents the cost of bringing dispersed parties together, negotiating terms, and revising agreements. We assume that all affected parties can participate when this contracting process takes place.

A public veto set P is imposed exogenously. It applies in every period and cannot be removed or waived by private agreement. The effective veto set within a contracting block is $P \cup S$: use occurs only if the initiator proceeds and every effective veto holder consents. The parties cannot purchase an exception to a public veto, even if they would jointly benefit from doing so. Political repeal is outside the model; P captures public authority that remains fixed over the horizon of private contracting. It is important to note that public vetoes play no positive welfare role in our model. This is assumed to isolate the costs public vetoes impose on naturally developing private order.

The same complexity that makes bargaining costly also limits informal cooperation between bargaining dates. The affected parties are dispersed and may have no continuing relationship with one another. We therefore rule out relational contracts over veto use: a future bargain may respond to information learned from use, but its rights and transfers cannot depend on who granted or withheld consent during the preceding block. This assumption gives R force. Parties can revise their formal agreements when they are brought together, but they cannot achieve the equivalent of continuous renegotiation through promises about future treatment.

A veto holder can nevertheless value what an experiment teaches it. With r periods remaining before the next bargain, let $v_{i,r}(s)$ denote party i's continuation value under the rights fixed for the current block. Its gain from permitting use is

$$D_{i,r}(s) = m_i + \delta \left(\mathbb{E}[v_{i,r-1}(s') \mid \text{use}] - v_{i,r-1}(s) \right), \tag{2}$$

where s' is the posterior after the payoff is observed. The first term is the expected current

payoff. The second is the effect of learning on the party's expected future payoff. A veto holder consents when $D_{i,r}(s) \geq 0$. If use is blocked, the belief state remains s. At each bargaining date, the parties choose S and transfers to maximize their joint expected value, subject to the continuing public veto.

4.3 Two costs of public ordering

The first cost is *suppressed learning*. If a public veto blocks use, the parties forgo the observation it would have produced. That observation might reveal a serious harm and lead them to abandon the application. It might instead reveal benefits, identify a manageable risk, or show which precaution works. By preventing the experiment, the veto can delay the information needed to develop more effective private agreements. This mechanism depends on learning through the activity being restricted. If the same information arrives independently, blocking use does not suppress learning through this channel.

The second cost is a *distorted allocation of veto power*. As information accumulates, the parties may discover that a different assignment of authority would permit uses whose total benefits exceed their costs. They can revise the private set S at the next bargaining date, using transfers to accommodate parties who would otherwise lose. They cannot revise the public set P through that bargain. A public veto can therefore continue to block an arrangement the parties would jointly choose, even after they have learned its likely consequences.

The two costs can interact. Blocking an experiment delays information that might support a new agreement, while the persistence of the public veto limits what the parties can do with information they acquire. Neither mechanism implies that every restriction reduces welfare. A veto may prevent an experiment whose expected harms exceed its benefits. The comparison depends on which uses the restriction blocks, what would have been learned from them, and which agreements the parties could have reached.

4.4 Open access within private order

At a bargaining date, the parties may assign private veto rights or choose $S = \varnothing$, leaving the initiator with open access during the coming block. A private veto protects parties

who expect losses from use. When $R > 1$, however, that assignment persists as information changes. A right that appeared sensible when assigned may prevent valuable experimentation before the parties can bargain again. Open access preserves those opportunities, but it can also permit unilateral uses whose harms exceed their benefits.

When $R = 1$, bargaining occurs before every opportunity to use the technology. Given full participation and transfers, the parties can choose private rights and agree on the current decision using their latest information. Open access has no distinctive advantage from avoiding an outdated private assignment. When $R > 1$, its flexibility may be valuable, but infrequent recontracting alone does not make open access preferable. The value of the additional experimentation must exceed the expected losses from uses that private veto rights would have prevented.

The principal distinction is between P and S. Both can block use today, but only S can be reassigned as the parties learn. Periodic contracting and the absence of stable relationships mean that private rights may also remain temporarily out of date. The model separates that temporary difficulty of revising private order from the more durable constraint imposed by a public veto. It provides a lens for understanding how governance of a frontier technology is discovered through use, and how restrictions on use can shape what private order becomes possible.

5 Discussion

AI governance requires institutions capable of responding to risks whose sources, magnitude, and incidence remain uncertain. Deployment can impose harms, but it also generates information needed to identify effective safeguards and determine who should exercise control. The cases and model suggest that public restrictions should be assessed partly by their effects on the discovery of private governance. An intervention may prevent a harmful application while also preventing the experience through which more discriminating institutions would otherwise develop.

The representation of private governance through veto rights captures one feature shared by otherwise different institutional arrangements. Certifiers can deny or withdraw approval,

affecting access to customers, insurance, and other commercial relationships. Insurers can withhold coverage from firms that fail to satisfy contractual requirements. Spectrum rights allow holders to exclude interfering uses. These mechanisms give particular parties conditional authority over activities involving their property, services, or reputations. Their influence can extend beyond the immediate contractual relationship when other market participants require insurance or certification as a condition of doing business.

These arrangements are richer than the binary consent decision in the model. Certification specifies precautions, insurance prices risk, and property rights permit negotiation over the terms of use. Withholding coverage or certification also need not prevent an activity altogether. The veto formulation abstracts from these differences to isolate how the allocation and duration of control affect experimentation. Its relevance depends on whether access to the resource or relationship being withheld is sufficiently important to constrain the activity. The model consequently represents a common mechanism within private governance rather than a complete account of its operation.

The cases illustrate how experience can change both the content of governance and the parties providing it. UL developed testing procedures, standards, and inspection practices as electrical products revealed unfamiliar hazards. Cyber insurers revise coverage, premiums, and security requirements as claims and technical developments change their understanding of risk. Early broadcasters negotiated interference disputes and developed claims to frequencies as commercial use revealed their value. In each case, governance required knowledge that was difficult to obtain before the technology was used. Experimentation supported the discovery of rules as well as the discovery of technological benefits and harms.

The model captures this process in a narrower form. Observed payoffs change beliefs about the consequences of use, which can change the preferred assignment of veto rights at the next bargaining date. It does not explicitly model the invention of safeguards. The cases establish the institutional plausibility of learning through use, while illustrating mechanisms beyond those formalized in the model. They should be understood as informing the theoretical comparison rather than directly testing its welfare implications.

The distinction between private and public authority concerns the means available for revising control. Private rights may become outdated between bargaining dates, and their

holders may block experimentation that would increase joint welfare. At the next bargain, however, affected parties can reassign those rights and negotiate transfers. The imposed public veto remains a constraint on that adjustment. Public authority can change in practice through legislation, but these processes differ from unanimous agreement among the affected parties.

The fixed public veto represents authority that persists over the relevant horizon of private contracting. This persistence produces two related costs. A veto can suppress learning by preventing the activity that would generate information. It can also prevent parties from implementing a mutually beneficial arrangement after information becomes available. The latter cost can remain even when uncertainty has substantially diminished: knowing that an application is valuable does not give private parties the authority to override a public restriction. The spectrum case illustrates how political intervention can displace emerging arrangements and substitute administrative decisions for subsequent bargaining (Hazlett 2020). UL and cyber insurance illustrate how standards and contractual controls can instead be revised through continuing institutional adjustment. Reversible use isolates these governance costs. An application can be stopped in the following period, so continued use is not technologically unavoidable. Nevertheless, an assignment of authority can persist after the information supporting it has changed.

The bargaining assumptions also limit the comparison. All affected parties can participate at contracting dates, realized payoffs are observable, and transfers are feasible. These conditions permit private agreements to account for interests that actual markets may leave unrepresented. Open access can also permit excessive experimentation when an initiator receives benefits while imposing larger losses on others. Its advantage is conditional on the value of additional experimentation exceeding these costs.

For AI, the implication is to preserve opportunities for institutional experimentation alongside technological experimentation. Even a framework employing competing private auditors may constrain this discovery process if public authorities fix the substantive standards those auditors must enforce. Public intervention may be warranted where expected harms are severe or private arrangements cannot represent affected interests. Its assessment should nevertheless include the information and governance arrangements that restrictions

may prevent from emerging.

6 Conclusion

AI governance is a problem of institutional discovery. The risks created by particular applications are difficult to know in advance, but so too are the institutions best suited to govern them. The historical cases illustrate this process. Experience with electrical hazards allowed Underwriters Laboratories to develop testing procedures, certification standards, labeling, and inspection. Cyber insurers use underwriting, claims, monitoring, and contractual requirements to revise governance as risks evolve. Early radio broadcasting revealed conflicts over interference and the value of particular frequencies, supporting the development of priority claims, contracting, and transferable rights. In each case, private order became more specific as experience accumulated.

The model isolates one mechanism within this broader process. Private veto rights can be revised as affected parties learn, while an imposed public veto cannot be removed through private agreement. This creates two costs of public ordering under uncertainty. A public veto can block the activity that would have generated useful information. It can also remain binding after learning makes a different allocation of authority mutually beneficial. The assumption of reversible deployment isolates these governance costs by ensuring that persistence comes from the institutional arrangement rather than from the technology itself.

Open access is not always preferable, nor can private ordering internalize every harm. Private rights may themselves remain temporarily out of date, and unilateral deployment can generate excessive experimentation. Private governance will also be weaker where affected parties cannot be identified or represented, harms are catastrophic, liability is limited, or agreements are difficult to enforce. The relevant comparison is among imperfect institutional alternatives rather than between an ideal market and an ideal regulator.

For AI, the implication is that public intervention should be evaluated by the information and governance arrangements it may prevent from emerging. Private actors possess contractual relationships through which increasingly specific controls can develop. Which mixture of exclusion rights will prove effective cannot be fully known in advance. The central lesson

is not that AI should remain ungoverned. It is that effective governance may itself be an outcome of experimentation. Public rules can sometimes improve that process, particularly where expected harms are severe or private arrangements cannot represent affected interests. But they can also suppress the learning and institutional adaptation needed to produce better governance. The challenge for AI policy is to address present harms without prematurely foreclosing the private order through which more effective controls may emerge.

References

[1] Alchian, A. A. (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy*, 58(3), 211–221.

[2] Arrow, K. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5), 941–973.

[3] Ball, D. W. (2025). A framework for the private governance of frontier artificial intelligence. *arXiv*. https://arxiv.org/abs/2504.11501.

[4] Beard, T. R., Ford, G. S., Spiwak, L. J., & Stern, M. (2014). Market mechanisms and the efficient use and management of scarce spectrum resources. *Federal Communications Law Journal*, 66(2), 263–308.

[5] Ben-Shahar, O., & Logue, K. D. (2012). Outsourcing regulation: How insurance reduces moral hazard. *Michigan Law Review*, 111(2), 197–248.

[6] Brearly, H. C. ([1923] 1997). A symbol of safety: The origins of Underwriters' Laboratories. In D. B. Klein (Ed.), *Reputation: Studies in the voluntary elicitation of good conduct* (pp. 75–84). University of Michigan Press.

[7] Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies "engines of growth"? *Journal of Econometrics*, 65(1), 83–108.

[8] Brough, W. T. (2025). More spectrum, please. *R Street Institute*. https://www.rstreet.org/commentary/more-spectrum-please/.

[9] Bustamante, P., Gomez, M., Murtazashvili, I., & Weiss, M. (2020). Spectrum anarchy: Why self-governance of the radio spectrum works better than we think. *Journal of Institutional Economics*, 16(6), 863–882.

[10] Coase, R. (1959). The Federal Communications Commission. *Journal of Law and Economics*, 2, 1–40.

[11] Coyne, C., & Leeson, P. T. (2005). Who's to protect cyberspace. *Journal of Law, Economics and Policy*, 1(1), 473–495.

[12] Dourado, E. (2012). Is there a cybersecurity market failure? Mercatus Center, George Mason University, Working Paper No. 12-05.

[13] Eling, M., Elvedi, M., & Falco, G. (2023). The economic impact of extreme cyber risk scenarios. *North American Actuarial Journal*, 27(3), 429–443.

[14] Geller, E. (2026). Fitch explains how water, healthcare organizations can keep strong credit ratings despite cyberattacks. *Cybersecurity Dive*, August 20, 2026. https://www.cybersecuritydive.com/news/water-healthcare-cyberattacks-credit-ratings-fitch/828384/.

[15] Greenblatt, R., Cotra, A., & Wijk, H. (2026). Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI/Hugging Face hacking incident. *METR*, August 26, 2026. https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/.

[16] Hadfield, G. K., & Clark, J. (2026). Regulatory markets: The future of AI governance. *Jurimetrics: The Journal of Law, Science and Technology*, 65, 195–240.

[17] Hazlett, T. (1990). The rationality of U.S. regulation of the broadcast spectrum. *Public Choice*, 33(1), 133–175.

[18] Hazlett, T. (2017). *The political spectrum: The tumultuous liberation of wireless technology, from Herbert Hoover to the smartphone*. Yale University Press.

[19] Hazlett, T. (2020). The 1927 Radio Act as pre-emption of common law property rights. *Review of Industrial Organization*, 56(1), 17–35.

[20] Hazlett, T., Palida, A., & Weiss, M. (2024). Governing complex externalities: Property rights for sharing radio spectrum. *Public Choice*, 201(3–4), 409–428.

[21] Herr, T. (2021). Cyber insurance and private governance: The enforcement power of markets. *Regulations & Governance*, 15(1), 98–114.

[22] *Hoover v. Intercity Radio Co., Inc.*, 286 F. 1003, 1007 (D.C. Cir. 1923).

[23] Hodgins, M. W. (2025a). The cybersecurity entrepreneur. *The Review of Austrian Economics*, 38, 265–286.

[24] Hodgins, M. W. (2025b). The perils of cybersecurity regulation. *The Review of Austrian Economics*, 38, 391–411.

[25] Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics*, 10(1), 74–91.

[26] Knowles, S. G. (2017). Private standards, public goods: Underwriters Laboratories and American safety engineering in the consumer era. *Historia Scientiarum*, 27(1), 42–62.

[27] Koppl, R. (2018). *Expert failure*. Cambridge University Press.

[28] Leeson, P. T. (2007). Efficient anarchy. *Public Choice*, 130(1–2), 41–53.

[29] Leeson, P. T. (2009). *The invisible hook: The hidden economics of pirates*. Princeton University Press.

[30] Lueck, D. (1995). The rule of first possession and the design of the law. *Journal of Law and Economics*, 38(2), 393–436.

[31] Lytton, T. D. (2014). Competitive third-party regulation: How private certification can overcome constraints that frustrate government regulation. *Theoretical Inquiries in Law*, 15(2), 539–572.

[32] Metcalf, T. (2026). AI safety and regulatory capture. *AI & Society*, 41, 2451–2466.

[33] Moore, T. (2010). The economics of cybersecurity: Principles and policy options. *International Journal of Critical Infrastructure*, 3(3), 103–117.

[34] Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.

[35] Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3), 641–672.

[36] Ostrom, V., Tiebout, C. M., & Warren, R. (1961). The organization of government in metropolitan areas: A theoretical inquiry. *American Political Science Review*, 55(4), 831–842.

[37] Paniagua, P., Rayamajhee, V., & Murtazashvili, I. (2024). Complex externalities: Introduction to the special issue. *Public Choice*, 201(3–4), 377–385.

[38] Powell, B. (2005). Is cybersecurity a public good? Evidence from the financial services industry. *Journal of Law, Economics and Policy*, 1(2), 497–510.

[39] Romanosky, S., & Robinson, C. (2026). *The insurability of artificial intelligence*. RAND Corporation, RR-A5130-1. https://www.rand.org/pubs/research_reports/RRA5130-1.html.

[40] Rowley, F. S. (1927). Problems in the law of radio communication. *University of Cincinnati Law Review*, 1(1), 1–35.

[41] Schumpeter, J. A. ([1934] 2021). *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle*. Transaction Books.

[42] Shortland, A. (2017). Governing kidnap for ransom: Lloyd's as a "private regime." *Governance*, 30(2), 283–299.

[43] Skarbek, D. (2011). Governance and prison gangs. *American Political Science Review*, 105(4), 702–716.

[44] Skarbek, D. (2014). *The social order of the underworld: How prison gangs govern the American penal system*. Oxford University Press.

[45] Sidak, J. G. (2016). Old regulations never die: Featherbedding and maritime safety after the Titanic. *Criterion Journal on Innovation*, 1, 201–227.

[46] Szpruch, L., Orfanoudaki, A., Maple, C., Wicker, M., Bengio, Y., Lam, K.-Y., & Detyniecki, M. (2026). Insuring AI: Incentivising safe and secure deployment. *Harvard Data Science Review*, 8(3).

[47] Tashima, S. (2026). Cyber insurance nears an inflection point. *Contingencies*, January/February 2026. American Academy of Actuaries. https://actuary.org/article/cyber-insurance-nears-an-inflection-point/.

[48] Thielman, N. (2026). Bootleggers, Baptists, and AI risk. *Fusion*, April 1, 2026. https://fusionaier.org/2026/bootleggers-baptists-and-ai-risk/.

[49] Thierer, A. (2023). *Flexible, pro-innovation governance strategies for artificial intelligence*. R Street Institute Policy Study No. 283. https://www.rstreet.org/research/flexible-pro-innovation-governance-strategies-for-artificial-intelligence/.

[50] Trout, C., Kvist, R., & Dattani, R. (2026). Bootstrapping frontier AI governance by mutualizing risk. *Lawfare*, September 22, 2026. https://www.lawfaremedia.org/article/bootstrapping-frontier-ai-governance-by-mutualizing-risk.

[51] *Tribune Co. v. Oak Leaves Broadcasting Station, Inc.*, Circuit Court of Cook County, Illinois, November 17, 1926. Reprinted in *Congressional Record*, 68 (1926), 216–219.

[52] *United States v. Zenith Radio Corp.*, 12 F.2d 614 (N.D. Ill. 1926).

[53] Viswanatha, A., & Krouse, S. (2025). Chinese spies hit more than 80 countries in "Salt Typhoon" breach, FBI reveals. *Wall Street Journal*, August 28,

2025. https://www.wsj.com/politics/national-security/chinese-spies-hit-more-than-80-countries-in-salt-typhoon-breach-fbi-reveals-59b2108f.

[54] Weil, G. (2026). Don't let AI developers hire their own referees. *AI Frontiers*, July 29, 2026. https://ai-frontiers.org/articles/dont-let-ai-developers-hire-their-own-referees.

[55] Wildavsky, A. (1988). *Searching for safety*. Routledge.

[56] Woods, D. W., & Wolff, J. (2025). A history of cyber risk transfer. *Journal of Cybersecurity*, 11(1), tyae028.