Preliminary Concept Paper

Designing Policymaking Mechanisms for Regulatory Dynamism

TTI REDESIGNING THE GOVERNANCE STACK
GEORGETOWN LAW

Preliminary Concept Paper

Designing Policymaking Mechanisms for Regulatory Dynamism

APRIL 2025

Julie E. Cohen
Nina-Simone Edwards
Meg Leta Jones
Paul Ohm

Co-Signers:[1]
Ellen Goodman
Ngozi Okidegbe
Jodi Short
Daniel Walters
Lauren Willis

---

1 The ideas in this concept paper were developed in part through workshops held with academics and regulators. Some of the attendees have decided, as well, to co-sign the paper. Although not every one of the co-signers agrees with every idea set forth below, each supports the proposals overall. In addition to the co-signers, we would like to thank Kendra Albert, Hilary Allen, William Boyd, Alissa Cooper, Laura Edelson, Casey Fiesler, Ellen Goodman, Laura Moy, Arvind Narayanan, Ngozi Okidegbe, Blake Reid, Neil Richards, Andrew Selbst, Jodi Short, Ranjit Singh, Alicia Solow-Niederman, Olivier Sylvain, Michael Veale, Daniel Walters, and Lauren Willis for their participation in the academic workshop and their constructive and generous feedback. We thank Britney Firmin for research assistance. This work was developed with the support of the Tech & Public Policy program at the McCourt School of Public Policy at Georgetown University, the Fritz Family Fellowship at Georgetown University, the Ford Foundation, and the Economic Security Project.

Executive Summary

This document contributes to a broader initiative focused on reimagining the role of the administrative state in the governance of a digital, data-driven economy. It examines needed changes to the design of regulatory policymaking mechanisms.

The administrative state is struggling to counter the harms of today's information economy. Existing mechanisms for policymaking fall short both substantively and procedurally. Substantively, regulators face challenges translating decisions about public values—e.g., "protect sensitive personal information" or "avoid deceiving consumers"—into forms capable of being operationalized within networked digital processes and environments. Procedurally, the regulatory toolkit is reactive and poorly adapted to iteration and experimentation, and the results it produces—sometimes, results that are already outdated—can be difficult to revisit as the information available to regulators and the public evolves.

In this report, we develop a set of foundational principles for the design of a regulatory system that is nimble and effective. These include: jumpstarting the regulatory lifecycle by empowering regulators to act sooner, enabling experimental approaches to regulation, creating governance seams to facilitate regulatory oversight, mandating beneficial friction at key points in networked digital systems and processes, and extending regulatory authority in ways that mirror the scale and interdependence of digital supply chains.

Next, we propose an expanded regulatory toolkit that implements these principles. To act in ways that effectively address digital architectures, systems, and processes, regulators must be empowered to mandate data flow restrictions, to develop design requirements for both user-facing and technical interfaces, to require continuous adversarial testing of certain kinds of systems and processes, and to develop and impose human subjects oversight requirements adapted to the operation of digital architectures, services, and supply chains.

Last, we propose corresponding institutional changes, including new statutory authorities to replace the relevant parts of the Administrative Procedure Act (APA) and their corresponding agency implementations. As a baseline, regulators must be empowered to engage in streamlined, iterative rulemaking and equipped with the resources to conduct interdisciplinary problem framing and assessment. Additionally, regulators should have authority to develop what we call policy sandboxes—experimental regimes of enhanced oversight that operate via tunable parameters—and to develop premarket certification and/or licensing regimes for digital architectures, products and services.

.

Introduction

The administrative state is struggling to counter the harms of today's information economy—including misinformation, pervasive state and commercial surveillance, and the growing entrenchment of data-driven bias and inequality. Existing policy mechanisms fall short both substantively and procedurally. Substantively, regulators face challenges translating decisions about public values—e.g., "protect sensitive personal information" or "avoid deceiving consumers"—into forms capable of being operationalized within networked digital processes and environments. Procedurally, the regulatory toolkit is reactive and poorly adapted to iteration and experimentation, and the results it produces—sometimes, results that are already outdated—can be difficult to revisit as the information available to regulators and the public evolves.

Networked digital systems and processes can create important private and social benefits. But they also can create significant private and social harms. We therefore conclude that networked digital architectures, systems, and processes will (and should) be subjected to regulatory oversight. The question is how. It is urgently important to design new policymaking mechanisms that enable regulatory experimentation while at the same time preserving appropriate public accountability. Further, such mechanisms should enable appropriately precautionary policymaking that addresses issues such as sensitive information inference, data and network insecurity, pernicious data-driven practices that entrench bias and economic precarity, and viral mis- and disinformation.

Part 1 provides an overview of four important factors that, taken together, have worked to disable regulators from making and effectively implementing policy choices for networked digital systems and processes. In Part 2, we develop a set of foundational principles for the design of a regulatory system that is nimble and effective: jumpstarting the regulatory lifecycle, enabling experimental approaches to regulation, creating governance seams to facilitate regulatory oversight, mandating beneficial friction at key points in networked digital systems and processes, and extending regulatory authority in ways that mirror the scale and interdependence of digital supply chains.

In Part 3, we present a suite of policymaking mechanisms designed to operationalize these principles. Through data flow restrictions, design requirements, continuous adversarial testing, and human subjects oversight, we show how regulators can reshape digital architectures, systems, and processes to better serve public values while preserving innovation and dynamism where they matter most. In Part 4, we outline corresponding institutional frameworks for implementing the mechanisms described in Part 3.

Three points are worth emphasizing at the outset: First and obviously, attaining regulatory dynamism also requires visibility into the processes and arrangements that are subject to regulatory oversight. The principles developed in Part 2 and the proposals in Parts 3 and 4 are designed to build on those in our earlier concept paper on regulatory monitoring,[2] which seek to introduce visibility where none currently exists. Second, as that concept paper also specifies, our recommendations are intended to widen the scope for regulatory experimentation only in the traditional domains of economic and social welfare regulation. In particular, very different oversight, due process, and rule-of-law considerations attach to law enforcement and border control activities, and we do not intend any recommendations about those activities.[3] Finally, redesign for regulatory dynamism might be pursued in various ways. The systems and institutions of concern to us affect the lives and livelihoods of hundreds of millions of people. Here and in this project generally, we are concerned with how to redesign regulatory institutions in ways that balance competing considerations of nimbleness, efficacy, public accountability, and the rule of law. None of our project recommendations in this document or elsewhere should be construed as endorsing approaches that do not prioritize all of these considerations.

2 Julie Cohen, et al., Regulatory Monitoring in the Information Economy (2024), https://www.law.georgetown.edu/tech-institute/wp-content/uploads/sites/42/2024/09/Regulatory-Monitoring-in-the-Information-Economy.pdf.

3 *See generally* Jessica Bulman-Pozen & Emily Chertoff, *The Administrative State's Two Faces*, LawFare (Feb. 24, 2025), https://www.lawfaremedia.org/article/the-administrative-state-s-two-faces.

Part 1: Regulatory Failure or Regulatory Mismatch?

The causes of regulatory failure in the networked information economy are many and interrelated. Some are ideological and cultural. As many capable scholars and advocates have shown, these include systematic misconceptions about what justifies regulation and systematic reinforcement of the idea that regulation is inevitably innovation-stifling. We are in substantial agreement with those critiques, which we sketch below, but fleshing them out in detail is not the primary focus of this report. Instead, we take them as our starting point and focus primarily on other problems that are procedural and institutional: Even if regulators were broadly empowered to innovate in pursuit of the public interest and wished to do so, they would be constrained on the front end by process paralysis and on the back end by the difficulty of translating the customary outputs of regulatory processes into forms that can be operationalized effectively within rapidly evolving networked digital environments.

What (and Who) Justifies Regulation

A central and much-remarked characteristic of the modern administrative state is the way it has reframed the project of regulatory oversight in a way that centers markets and grudgingly permits regulators to correct certain types of market failure so long as they do not impose unacceptably high costs on market participants. This priority manifests most obviously in the cost-benefit analysis of proposed agency rulemaking mandated by the Office of Information and Regulatory Affairs (OIRA); in the budgeting practices imposed on agencies by the Office of Management and Budget (OMB); and in the form of Bureaus of Economics at agencies such as the FTC and FCC.[4] But it also manifests in many other ways: in the evidentiary burdens that regulators must meet to justify their proposed actions; in the processes described below that leave regulators continually racing to catch up to evolving private-sector business models, technologies, and practices; in the bipartisan lack of interest (albeit for very different reasons) in more drastic and systematic redesign for regulatory empowerment; and even in the language used to describe regulatory actions (typically, as "burdens" or as "interventions" in otherwise-smoothly running processes).

The framing of markets and their outputs as virtuous and regulatory oversight inevitably both burdensome and unproductive is wrong both as a matter of history and as a matter of logic. Regulatory oversight and economic growth go hand and hand, and both are necessary to build a strong and resilient economy. So, for example, at the dawn of the industrial era, technologies for industrial production initially developed without regulatory oversight, generating both significant private and social economic benefits and significant private and social harms. Regulatory oversight worked to rebalance the scales, asserting the public interest in matters such as worker safety, worker rights, consumer protection, and environmental protection and subjecting industrial producers to new sets of obligations.[5] During the first half

4 *See* K. Sabeel Rahman, *Modernizing Regulatory Review*, Reg. Rev. (May 15, 2023), https://www.theregreview.org/2023/05/15/rahman-modernizing-regulatory-review/; Eloise Pasachoff, *The President's Budget as a Source of Agency Policy Control*, 25 Yale L. J. 2182 (2016).

5 *See, e.g.,* Thomas O. McGarity, Freedom to Harm: The Lasting Legacy of the Laissez Faire Revival (2013); Leonard Kennedy, Patricia McCoy, & Ethan Bernstein, *The Consumer Financial Protection Bureau: Financial Regulation for the Twenty-First Century*, 97 Cornell L. Rev. 1141 (2012).

of the twentieth century, regulatory coordination mobilized market participation in the construction of large infrastructures, such as highways and electrical grids, that served both private and social needs as they were then conceptualized.[6] For exactly these reasons, of course, we do not claim that regulatory oversight always works well; if it did, there would be no need to undertake this project. Regulatory mechanisms and institutions designed for the problems of an earlier era can fail to translate well to new ones. In such cases, a comprehensive redesign may be needed.

More generally, a framing of the regulatory mission that centers markets and foregrounds the costs of regulation ignores that cost and benefits are often in the eye of the beholder–what seems best for the firm or industry under scrutiny works by creating costs and attendant inefficiencies for others. Today, for example, so-called "dark patterns" in user interfaces work to prevent users from making choices inconsistent with the business models of system deployers. Such interfaces work well in a cost-benefit sense for the businesses involved, but they impose costs on users and also undermine broader societal goals such as, for example, fostering trust in consumer markets and narrowing the attack surface for identity theft. As another example, delayed trading mechanisms designed to prevent some forms of high-frequency trading restrict one kind of opportunity for profit extraction, but they also limit market volatility and foster increased trust in trading systems and intermediaries.

We begin from the premise that regulators should be empowered to regulate the technologies that now shape and permeate everyday life in the service of important public values and priorities.

Who Innovates, and Why

Related to the too-narrow focus on market benefits and regulatory burdens is a commitment to clearing the way for private-sector innovation. In particular, the networked digital economy and the architectures and processes that comprise it are said to permit continual innovation, which regulatory oversight would stifle. In this framing, all other public values are less important than progress—even when the destination is unclear or even foreseeably harmful. Also in this framing, regulators themselves do not, cannot, and should not attempt to innovate–even as private-sector innovators are encouraged to try and fail repeatedly.[7] This produces a culture of regulatory defeatism and enfeeblement, in which regulators doubt their own institutional capacity for positive change.[8]

The framing of regulatory oversight as innovation-stifling is wrong both as a matter of history and as a matter of logic. To begin, it ignores the historical roles of government investment and publicly determined priorities in shaping significant innovations of the modern era, including the internet, nuclear power, and air travel.[9] Additionally, it ignores the interplay between private innovation and background social and public values and priorities. So, for example, as motor vehicle technology evolved, the public interest in first automotive safety and later environmental sustainability has encouraged research on passive restraints and alternative fuels. As pharmaceuticals became more sophisticated, public regulation has evolved to emphasize the need for controlled testing as a mechanism for assessing efficacy and documenting adverse effects. Food safety rules and monitoring requirements have protected public health

6 *See* William Boyd, *Decommodifying Electricity*, 97 So. Cal L. Rev. 937, 1012-13 (2024); Jacob Hacker & Paul Pierson, American Amnesia: How the War on Government Led Us to Forget What Made America Prosper (2016).

7 *See, e.g.,* Kenneth Chang, *Twin Test Flight Explosions Show SpaceX Is No Longer Defying Gravity*, N.Y. Times (Mar. 8, 2025), https://www.nytimes.com/2025/03/08/science/starship-spacex-explosion-elon-musk.html (describing the repeated explosions of SpaceX rockets as "not necessarily failures for a company that has thrived on a mind-set of 'launch it, break it, fix it, launch again.'").

8 *See* Jodi Short, *Regulatory Managerialism as Gaslighting Government*, 86 L. & Contemp. Probs.1 (2023).

9 *See* Janet Abbate, *Government, Business, and the Making of the Internet,* 75 Bus. Hist. Rev. 147 (2001); Mariana Mazzucato, *Mission-Oriented Innovation Policies: Challenges and Opportunities*, 27 Indus. & Corp. Change 803 (2018); H.W. Lawrence, Aviation and the Role of Government (2004).

and limited the spread of pathogen outbreaks. And in each of those fields, regulatory oversight and, importantly, regulatory innovation have played important roles in reshaping the trajectory of private innovation to align more closely with changing public priorities.

We begin from the premise that regulators should be empowered to innovate in the service of important public values and priorities.

What Causes Process Paralysis

Today's regulatory processes are slow and barnacle-encrusted, taking years to complete and lending themselves readily to stalemate and capture. Some of this is by design: the formal rulemaking process described in the Administrative Procedure Act envisions multiple, distinct stages, including providing notice of the proposed rulemaking, appointing an administrative law judge, collecting evidence, and conducting a hearing which includes cross-examination. But even the so-called "informal" rulemaking process[10] that has become the norm and that was initially intended to work somewhat more speedily now takes, on average, 4 years.[11]

We do not mean to suggest that regulatory processes should be perfunctory or rushed. Attaining complex regulatory goals requires both speed and rigor. The problem is the way process paralysis can result from organized resistance by regulated industries and entities (or those wishing to avoid regulation). In particular, well-resourced actors know how to weaponize opportunities for public comment, whether by mounting organized "astroturfing" campaigns designed to create a misleading perception of strong public opposition to a proposed rule or by using comments to lay the groundwork for litigation designed to prevent a rule from taking effect.[12] Even when they aren't acting so maliciously, contributions from industry stakeholders and other well-resourced actors tend to receive disproportionate attention from regulators, and such contributions are far more likely to counsel caution than they are to recommend regulatory innovation.[13] Last but not least, well-intentioned reforms have imposed successive layers of process requirements and restrictions, magnifying the public perception of government as a vast, unresponsive, wasteful bureaucracy.[14]

This state of affairs reinforces the ideological and cultural barriers to regulation described above in two ways. First, it contributes significantly to what has become known as the "pacing problem"—the idea that regulators can only react to new technologies and, therefore, are limited in what they can hope to achieve. Legacy process paralysis diminishes the efficacy of regulatory actions both by hindering regulators from responding quickly and effectively to emergent risks and by frustrating efforts to undertake proactive oversight. Second, process paralysis fuels popular anger at the state's seeming inability to act in the

---

10 Administrative Procedure Act, 5 U.S.C. § 553 (1966).

11 A GAO study of 139 major rules found that "the average time needed to complete a rulemaking across our 16 case-study rules was about 4 years, with a range from about 1 year to nearly 14 years, but there was considerable variation among agencies and rules." Government Accountability Office, Improvements Needed to Monitoring and Evaluation of Rules Development as Well as to the Transparency of OMB Regulatory Reviews (2009); *see also* Curtis Copeland, Cong. Rsch. Serv., RL 32240, The Federal Rulemaking Process: An Overview 3 (2013).

12 *See Ohio v. EPA*, 603 U.S. 279 (2024) (holding that failure to consider large volume of comments coordinated by trade associations, states, and several companies rendered the EPA's revised air quality standards arbitrary and capricious); Edward Walker, *Millions of Fake Commenters Asked the FCC to End Net Neutrality. 'Astroturfing' Is a Business Model*, Wash. Post (May 14, 2021), https://www.washingtonpost.com/politics/2021/05/14/millions-fake-commenters-asked-fcc-end-net-neutrality-astroturfing-is-business-model/ (describing comment campaign secretly orchestrated by broadband companies).

13 *See* Sidney Shapiro, Elizabeth Fisher & Wendy Wagner, *The Enlightenment of Administrative Law: Looking Inside the Agency for Legitimacy*, 47 Wake Forest L. Rev. 463, 464 (2012); Jason Webb Yackee & Susan Webb Yackee, *A Bias Towards Business? Assessing Interest Group Influence on the US. Bureaucracy*, 68 J. Pol. 128 (2006).

14 *See* William Funk, *Public Participation and Transparency in Administrative Law-Three Examples As an Object Lesson*, 61 Admin. L. Rev.171 (2009); Stuart Shapiro, *The Paperwork Reduction Act: Benefits, Costs and Directions for Refo*rm, 30 Gov't Info. Q. 204, 206 (2013); Adam Raviv, *Government Ethics in the Age of Trump*, 54 U. Mich. J. L. Reform 331 (2021).

public interest. Such anger may be perceived–we believe inaccurately–as lending cover to efforts to dismantle the administrative state or some of its functions, effectuating a retrenchment that would prevent it from acting in the public interest at all. Our project here is different.

We begin from the premise that regulators should be empowered to exercise proactive, nimble oversight of regulated industries and entities.

Institutional Origins of Operational Detachment

In the existing regulatory landscape, implementation problems arising in the vast and widening "middle space" between rulemaking and enforcement have given rise to a large and varied set of implementation problems.

First, broad public mandates addressed to entire industries–e.g., "don't discriminate based on protected characteristics," "protect sensitive personal information," or "avoid deceiving consumers"–need to be translated into forms that can be operationalized in networked digital environments. As a practical matter, then, the actors in networked information industries may be left with significant operational discretion regarding how to implement public mandates. In a sense, this has been true for regulation of industrial processes, but networked digital ecosystems offer flexible, interconnected products and services. The interactions between such products and services are many and varied. As a result, it is more difficult both to identify resulting harms and to assess the extent of effective compliance with public mandates.

Additionally, networked digital architectures, systems, and processes are very often configurable and reconfigurable in real time. Systemic workarounds that counteract the effects of public mandates can arise from the ordinary operation of machine learning processes that seek to optimize for certain outcomes while routing around obstacles to those outcomes. This type of flexibility significantly undercuts the efficacy of traditional forms of regulatory policymaking.

These problems, in turn, have engendered sets of processes, tools, and actors focused on questions of compliance. Yet regulators themselves often lack the methods and resources to oversee compliance directly. Compliance monitoring and assessment functions therefore are outsourced to an array of third-party intermediaries providing a variety of different services.[15] Such intermediaries have their own incentives, however, and ever more complex compliance ecosystems do not necessarily produce improved results. At worst, intermediary involvement can distort the regulatory field of vision. More commonly, it can result in a performative, checkbox approach to determining whether public mandates have been satisfied.[16]

From one perspective, these are enforcement problems, which we consider in a future module of this project. From another, however, they are problems with the institutional structure for public policymaking, which is insufficiently attentive to the growing need to involve policymakers more directly in questions about how to operationalize public mandates.

We begin from the premises that regulators should be empowered to narrow the gap between broad policy mandates and enforcement by crafting operational mandates for entities under their jurisdiction and that this power should include the authority to craft and experiment with operational mandates specific to particular entities.

---

15 *See* Kenneth A. Bamberger, *Technologies of Compliance: Risk and Regulation in a Digital Age*, 88 Tex. L. Rev. 660, 712 (2010).

16 *See, e.g.,* Ari Ezra Waldman, *Privacy, Practice, and Performance*, 110 Cal. L. Rev. 1221, 1239 (2022); Lauren Edelman, *Legal Ambiguity and the Politics of Compliance: The Case of Sexual Harassment Policies in the Workplace*, 27 Law & Soc'y Rev. 663 (1993).

Part 2: Principles of Regulatory Dynamism

A different and, we think, more productive way to think about the problems cataloged in Part 1 is that industrial and regulatory processes are mismatched in five important ways, which relate to timing, operational immediacy, leverage, pace, and scale. If proposals for redesign of regulation are to succeed, they need to build on the principles described here.

Timing: Jumpstarting the Regulatory Lifecycle

The first mismatch between industrial and regulatory processes relates to timing. Existing mechanisms for regulatory policymaking force regulators to move slowly both by design and because they are enfeebled by a culture of regulatory disempowerment. Meanwhile, the need for regulatory dynamism has only grown more acute. Tethered and/or platformized services can be altered and tuned in real time. If regulatory oversight is to be effective, it requires a correspondingly nimble lifecycle.

This observation about timing mismatch is not new. From time to time, scholars have proposed adjustments intended to make regulatory processes nimbler. One instructive example is the call for adaptive management frameworks within environmental regulatory processes.[17] These proposals emphasized the importance of monitoring, assessment, feedback, and iterative adjustment. Yet (at least in the U.S.) neither legislators nor, to the extent feasible, regulators themselves turned those aspirations into concrete process mechanisms. Instead, regulators have sometimes engaged in what scholars called "a/m lite," committing to "learning while doing" without adopting any specific method or process for doing so.[18]

A different kind of example is the shift to performance-based regulation in a number of regulatory fields including consumer financial protection and public utility regulation.[19] Under performance-based regulation, firms themselves must experiment with ways of meeting designated targets that regulators specify or eliminating harms that regulators have identified. Such approaches can be extremely effective when the universe of regulated firms and products is well defined and the performance standard is easy to specify; they are less effective when regulated firms and/or their activities are more varied, when the activities are more difficult to monitor, and when other workarounds can offset the formal attainment of designated performance levels.[20]

In short, regulatory nimbleness can't simply be melded into today's regulatory mechanisms while leaving those mechanisms otherwise unaltered. Attaining regulatory nimbleness requires a fundamentally different approach to structuring regulatory authority.

17 *See* Lawrence E. Susskind & Joshua Secunda, *"Improving" Project XL: Helping Adaptive Management to Work Within EPA*, 17 UCLA J. Envtl. L. & Pol'y 155, 166-67 (1999); Lawrence Susskind et al., *Collaborative Planning and Adaptive Management in Glen Canyon: A Cautionary Tale*, 35 Colum. J. Envtl. L. 1, 4 (2010); Justin R. Pidot, *Governance and Uncertainty*, 37 Cardozo L. Rev. 113 (2015); Jeremy A. Carp, *Autonomous Vehicles: Problems and Principles for Future Regulation*, 4 U. Pa. J. L. & Pub. Aff. 81 (2018).

18 *See* Robin Kundis Craig & J.B. Ruhl, *Designing Administrative Law for Adaptive Management,* 67 Vand. L. Rev. 1, 14 (2014); *see also* Jason Thomann, et al., *Adaptive Management in Groundwater Planning and Development: A Review of Theory and Applications*, 586 J. Hydrology 1 (2020) (suggesting that adaptive management techniques may be poorly suited to managing environmental impacts, particularly where impacts are irreversible).

19 *See* Lauren Willis, *Performance-Based Consumer and Investor Protection: Corporate Responsibility Without Blame, in* The Culpable Corporate Mind (2022); Emma Shumway, et al., *Addressing Energy Insecurity Upstream: Electric Utility Ratemaking and Rate Design As Levers for Change*, 45 Energy L.J. 361, 374 (2024).

20 *See generally* Cary Coglianese, *The Limits of Performance-Based Regulation*, 50 U. Mich. J. L. Reform 525, 545-47 (2017).

Operational Immediacy: Enabling Regulatory Experimentation

To respond to the rapid, iterative and interactive lifecycle of today's networked digital products and services, the regulatory lifecycle also needs to be similarly iterative and interactive, enabling dynamic and collaborative policy implementation and tuning. Today's processes put too much weight and pressure on getting regulatory action precisely tailored from inception. Even the nomenclature of a two-step "proposed rule" and "final rule" betrays an unrealistic belief that two steps and a few months (or, often, years) are all that are required to get an agency action right. Dynamic industrial processes require correspondingly dynamic responses.

This observation about the need for dynamism in regulatory oversight also is not new. From time to time, scholars have called on government to adopt a culture of regulatory experimentalism.[21] As with regulatory nimbleness, however, regulatory experimentalism can't be attained if existing policymaking mechanisms are left unaltered. Attaining regulatory nimbleness requires different mechanisms and a different approach to structuring the regulatory lifecycle. Giving agencies the freedom to experiment means authorizing them to act earlier and under more uncertainty, giving them the time and tools they need to measure and monitor industry responses, and empowering them to tune and update regulatory mandates and guardrails.[22]

It starts at the beginning. Agencies must be empowered to spend less time and energy crafting final agency actions and, instead, to implement and experiment with preliminary actions under leaner, less burdensome and less bureaucratic processes. Where such authority currently exists, it has been conceptualized as deregulatory.

As we explain in Part 4 below, we propose to "flip the script," empowering agencies to test new requirements for the design and/or operation of networked digital services and processes.

To reap the full benefit of experimentation and iterative improvement, such early regulatory actions should be designed to elicit observable responses from regulated entities and other relevant stakeholders. In some cases, the effects of an experimental policy regimen might be observed and measured by directly interacting with a company's website or app or by examining changes to its terms of service. In others, the experimental regimen might include mandatory reporting or monitoring obligations.

To fuel experimentation, regulatory mandates should be devised, when possible, to be tunable. For example, an experimental intervention designed to tamp down on harmful social media virality might limit the number of times a message can be reposted. The numeric parameter could be increased or decreased over time, enabling the agency to adjust to what it is learning, rather than forcing it to start over from scratch with a formal process to amend a fixed rule.

For this approach to work, post hoc judicial review should also build in leeway for agency experimentation. This observation implicates recent developments in the courts that are outside the scope of this concept paper. A later module of this project will be devoted to the rule of law requirements that should attach to the operation of a redesigned, nimbler administrative state.

---

21 *See* Charles F. Sabel & William H. Simon, *Minimalism and Experimentalism in the Administrative State*, 100 Geo. L. J. 53 (2011).

22 In other words, it requires expressly empowering agencies to experiment (and thereby overruling case law to the contrary). *See, e.g., New York Stock Exch. LLC v. Sec. & Exch. Comm'n*, 962 F.3d 541 (D.C. Cir. 2020) (vacating an SEC rule establishing an experimental program to test fee structures on cost-benefit grounds).

Leverage: Creating Governance Seams

The third mismatch between industrial and regulatory processes relates to the industrial preference for seamless continuity in networked and/or platformized processes. Private actors celebrate seamlessness both in the "user experience" and in the way they structure the back ends of their businesses. Seams and separations are criticized as intrinsically inefficient. But seamlessness can create problems for users, competitors, and regulators. For example, seamless integration of multiple services can lock users into platforms and ecosystems, making it difficult or impossible for independent providers to compete with services offered by dominant platforms and difficult or impossible for users to disable interconnected services that they do not want. Fully integrated and seamless digital architectures can also obscure information flows in ways that make it difficult for regulators to exercise effective and public-regarding oversight.

By contrast, seams and discontinuities in technological systems can create valuable points of visibility and intervention, making system operations more legible and creating natural break points for understanding–and imposing–governance.[23] Seams and separations therefore should not automatically be viewed as inefficiencies to be eliminated. Thoughtfully designed, they are features, not bugs, and they can be intentionally preserved or created to enable oversight and accountability.

Regulators should be empowered to mandate governance seams within digital architectures, systems, and processes, both to address specific problems and to preserve the ability to protect articulated public values at the systemic level.[24] For instance, a regulator wanting to facilitate third-party competition with dominant platform services might require the dominant platform to create and maintain the undiminished functionality of interfaces between its services that could serve as interconnection points. Such interfaces also might become sites of transparency. Just as a border in the geographic physical world gives a place to inspect the goods flowing into a nation, so too are seams between software services a place to position a regulatory inspector. Well-placed and well-designed governance seams can serve both as points of access and points of oversight.

Pace: Friction and Efficiency in Balance

A fourth mismatch between industrial and regulatory processes relates to the speed of networked digital processes. As a needed corrective to the singular focus on the efficiency of business operations, regulators should be empowered to introduce beneficial friction into digital architectures, products, and services.[25] Well-placed constraints or delays can create time and space for human judgment, reflection, transparency, or intervention, ensuring that digital systems operate at speeds compatible with human needs and public values.

Regulators might inject friction directly: for example, an agency might promulgate a rule or negotiate a consent decree requiring a social media platform to replace the "infinite scroll" style of presentation with mandatory page breaks. They might impose friction indirectly, by affecting the incentives that push companies toward seamlessness and scale: for example, an agency could impose stricter obligations on companies that permit users to communicate with millions of other users than on companies that use design to cap message group size to foster community building and self-regulation.

---

23 Cf. Janet Vertesi, *Seamful Spaces: Heterogeneous Infrastructures in Interaction*, 39 Sci. Tech. & Hum. Values 264 (2014); Julie E. Cohen, Configuring the Networked Self: Law, Code, and the Play of Everyday Practice 239 (2012).

24 *See* Brett Frischmann & Paul Ohm, *Governance Seams*, 37 Harv. L. J. & Tech. 1117 (2023).

25 *See* Brett Frischmann & Susan Benesch, *Friction-in-design Regulation as 21st Century Time, Place, and Manner Restriction*, 25 Yale. L. J. & Tech. 376 (2023).

The point is not to replace efficiency with friction as an all-encompassing end goal for society or the economy. Regulators should seek to create the conditions for efficiency where it is needed and the same for friction. The point is to use the appropriate means (efficiency or friction) to achieve the stated goals of regulation.

Scale and Scope: Extending Regulatory Authority Along Digital Supply Chains

A final mismatch between industrial and regulatory processes relates to scale and interdependence. The digital economy is networked and interconnected, with many smaller actors nested within larger ecosystems and digital supply chains. Even when particular actors retain control over some component(s) of their own systems and services, they may lack good information about the provenance of data, algorithms, and models flowing into those systems and services and/or about the destinations and uses of information flowing out. Regulators operating with constrained visibility and/or limited jurisdiction are not much better off. Unaccountable data and algorithmic supply chains threaten effective public oversight of digital architectures, systems, and processes–and enable industry to practice novel forms of regulatory arbitrage. It has become important to think of the regulatory process as itself involving a supply chain of sorts: a dynamic framework for comprehensive and effective oversight of networked, interconnected systems and processes that matches the digital supply chains for such processes in scale and scope.

This observation about scale mismatch also is not new. For several decades now, scholars in fields such as financial regulation, environmental regulation, and energy regulation have underscored the need to exercise systemic oversight of interconnected processes. Regulatory processes in those fields have evolved correspondingly. By contrast, although scholars who study the digital economy have urged more comprehensive attention to the scale and complexity of modern digital supply chains, regulatory oversight of networked digital architectures, systems, and processes has been relatively inattentive to such matters, and that has remained the case even as digital supply chains have woven different sectors of the economy ever more tightly together. That needs to change.

We note, finally, that there are two questions related to this recommendation that this concept paper does not address. One concerns the executive branch org chart. Most basically, agencies charged with oversight of digital economy processes should be empowered to do their jobs, but responding effectively to the challenges of networked digital architectures, systems, and processes may require some restructuring. The Digital Architectures, Services, and Processes Oversight Board described in our concept paper on regulatory monitoring could serve as a fulcrum point for that endeavor. Effective restructuring also may require creation of more streamlined and effective processes for interagency communication and/or coordination. We plan to take up these issues in a later stage of this project.

A second question that sits largely outside the scope of this concept paper relates to the fact that digital and informational supply chains are global. Crafting delegations of regulatory authority that respond appropriately to the resulting challenges is a task for Congress. One strategy might involve empowering regulators to protect U.S. residents by crafting mandates with extraterritorial bite. So, for example, the European General Data Protection Regulation binds firms processing data of European residents even if they are not physically present in Europe. We acknowledge, as well, that thinking on the extent to which national security considerations should inform regulatory oversight of networked digital architectures, systems, and processes is still evolving.

In Part 3, we present a suite of policy mechanisms designed to operationalize these principles.

Part 3: Elements of the Dynamic Regulatory Toolkit

As is by now well documented, current disaggregated networked digital architectures, systems, and processes are designed to collect large volumes of data and process it with the goal of optimizing outputs in real time for particular purposes. Such arrangements have proved highly resistant to intervention at the endpoints for reasons that relate to considerations of timing, operational immediacy, leverage, pace, and scale. Operationalizing public mandates requires an expanded regulatory toolkit. The mechanisms described in this Part are intended to realign regulatory and industrial timescales, enable regulatory experimentation, facilitate the creation of governance seams, allow regulators to introduce beneficial friction into networked digital processes, and generate regulatory actions matched to the scale and complexity of digital ecosystems and supply chains.

At the outset, it is worth underscoring the shortcomings of two kinds of approaches on which proposals for regulatory reform commonly rely. One common regulatory frame, familiar from longstanding debates in fields as varied as consumer protection, commercial contracting, privacy, investor protection, and at-will employment, is notice and consent. Machine learning based digital processes, however, operate in real time and at scale in ways that are difficult to understand.[26] Offloading the need for repeated consent onto users of such processes (including not only individuals but also downstream business users) imposes exhaustion and attention overload and requires wasteful and duplicative allocation of organizational resources.[27] And, because the consent of a relatively small proportion of users is often sufficient to train and optimize digital systems and processes in ways that affect all users, the control afforded even by very frequent requests for consent is often illusory. So too with endpoint based restrictions that focus on downstream use cases or on subsequent sale or transfer of particular kinds of information after it is collected. Such restrictions come too late, do too little, and don't focus on the workarounds that scaled-up processing can enable. Both consent and data use restrictions have continuing roles to play in the context of a more comprehensive regulatory toolkit, but they cannot do all of the work that has been asked of them.

Instead, like the principles in Part 2, our recommendations here are informed by private-sector practices. It is clear that parties who design and deploy networked, data-driven, machine learning based processes have developed methods for designing, evaluating, and tuning them. The design of a dynamic toolkit for regulators should borrow from those methods.

Data Flow Restrictions

Data flow restrictions provide regulators with specific tools for introducing beneficial friction and governance seams into digital architectures at multiple scales. These restrictions can operate both at user-facing interfaces and deep within platform infrastructure, creating carefully calibrated points of intervention that slow down, separate, and make visible otherwise seamless data processes. By enabling regulators to modulate the speed and scope of data collection, storage, and use, these mechanisms help address fundamental mismatches between regulatory and industrial time scales.

26 *See* Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, 3 Big Data & Society 1 (2016).

27 Cf. Kiel Brennan-Marquez & Brendan Maher, *Micro Costs*, __ Geo. L. J. (forthcoming 2025).

At the user interface level, data flow restrictions can directly shape user experience in ways that promote reflection and agency. For example, replacing infinite scroll mechanisms with page breaks introduces moments of pause in social media consumption, while rate limits on message forwarding can help stem the viral spread of misinformation without engaging in content-specific regulation.

Other kinds of data flow restrictions can operate internally to digital architectures, systems, and processes. So, for example, the regulator could set and iteratively adjust parameters that might include: restrictions governing data collection (e.g., what qualifies as prohibited data, because it measures attention and/or browsing behaviors) and/or data storage (e.g., can some data be collected as long as it is stored, securely, only in-app, only on-device, etc.) and/or use categories (e.g., that are white-listed or black-listed or handled some other way).

This kind of intervention might apply to particular categories of sensitive information: a rule might focus on geolocation information, say, mandating that latitude/longitude information revealing the location and movement of users must be stored securely on-device and accessed only for particular purposes or by a registered list of apps, or even forbidding it from being collected at all.

Another example of such an approach might involve the software development kits (SDKs) furnished to app developers by platform proprietors. SDKs make application development faster and easier while simultaneously enabling platform controls and platform data collection practices. By targeting SDK requirements, regulators can shape how data flows are structured across entire ecosystems of third-party applications, creating governance seams that make data practices more transparent and accountable. SDKs therefore represent an important and overlooked source of leverage for regulators wanting to define and ensure implementation of public mandates.

Data flow restrictions allow for the benefits of tunability described above. Parameters like the number of posts before a pause, the threshold for viral message forwarding, or the rate limits for data access can be adjusted based on evidence about their effectiveness. This enables an iterative approach where regulators can start with initial estimates and refine restrictions based on observed outcomes.

Importantly, data flow restrictions can often address content-related challenges through content-neutral means. Rather than attempting to directly regulate specific types of harmful content, carefully designed friction in data flows can help mitigate broader systemic risks like viral misinformation or privacy violations. This approach may prove more resilient to First Amendment challenges while effectively serving public interest goals.

Design Requirements

Design requirements offer regulators another mechanism for addressing both the temporal and scalar mismatches that characterize current regulatory approaches to digital interfaces and the ways they frame access to underlying products, services, and architectures. Rather than waiting to react to problematic design patterns and practices after they emerge and cause harm, regulators should be empowered to establish and mandate adherence to clear specifications for the design of digital interfaces from the outset.

Unlike the recommendations that now appear in agency best practices statements and other guidance statements, the "design requirements" we envision are specific criteria that the designers of digital interfaces and architectures must meet. In this respect they are more similar to common technical standards that might be created with coordination by a standards body and subsequently serve as a commonly agreed basis for technical functionality. Unlike such standards (which we will address in a future module of this project), design requirements are specifications created under the direction of an agency and subject to ongoing agency supervision and iteration.

Consider how the Federal Trade Commission has historically approached deceptive interface design. The agency has repeatedly challenged companies that create dark patterns—interfaces designed to manipulate users through difficult-to-cancel subscriptions, disguised advertisements, buried terms, and deceptive data collection practices. Yet this case-by-case enforcement approach struggles to keep pace with the rapid evolution of interface design and fails to address the systemic nature of these practices.

Instead of continually reacting to new variants of deceptive design, the FTC could develop comprehensive design requirements that establish clear benchmarks for "light patterns" that respect user autonomy and facilitate genuine choice. The agency would work with an interdisciplinary team of experts in human-computer interaction (HCI) and related subjects to develop mandatory rules addressing matters such as color arrangements, length of text disclosures and prompts, number of tolerable steps or clicks, conventions for disclosing prices and other costs (including data) of services, and how to present errors or changes in status. The HCI team should engage with research communities and conduct independent research to establish the requirements. To facilitate their implementation, design requirements could be tested on a small scale before mandating their adoption. Regular review cycles could enable updates based on emerging research, changing user needs, and new technological capabilities.

Design requirements are not for user interface design only; they can dictate the design of "back end" interfaces and architectures not visible to end users. Design requirements like these are particularly well-suited to address the interconnected, nested nature of modern digital ecosystems. They can be structured to create consistent criteria that propagate throughout digital supply chains, scale horizontally across different types of interfaces and vertically through technology stacks, account for how interface elements interact across connected services and platforms, and provide clear instruction to both large platforms and smaller actors operating within platform ecosystems. These scalar advantages become especially apparent in areas like component libraries that get reused across multiple products, SDK interfaces that mediate between platforms and third-party developers, authentication flows that need to work consistently across connected services, and data collection interfaces that affect multiple downstream data processors.

Continuous Adversarial Testing Requirements

For some types of digital architectures, products, and services, requirements for continuous adversarial testing will be necessary to enable regulators to identify and address potential harms before they manifest at scale. Unlike traditional testing approaches that focus on specific, predefined harms, adversarial testing creates structured opportunities to discover unanticipated problems and assess them within the broader context of interconnected digital systems.[28] Requirements for continuous adversarial testing create crucial governance seams in the technology deployment lifecycle.

The need for this approach is particularly acute for complex technologies that resist straightforward interpretation or explanation—for example, many systems built on machine learning models. While agencies like the FDA and Federal Reserve have experience testing relatively simple technologies for specific, well-defined harms, modern digital systems present a broader and more dynamic set of risks. Some harms may not emerge until well after deployment, while others may arise from complex interactions between multiple systems and distributed data architectures that the original developer neither controls nor can fully test.

To implement continuous adversarial testing effectively, companies must notify regulators about systems they intend to launch, deploy in new ways, or significantly adjust, with sufficient lead time for thorough evaluation. Government experts—including approved researchers from civil society organizations and academia—can then work with agency staff to conduct comprehensive testing regimens. These assessments should range from standardized test suites to human-driven experimentation designed to uncover potential failure modes and negative impacts.

28 Such testing already plays an important role in some regulatory fields. *See generally* Mehrsa Baradaran, *Regulation by Hypothetical*, 67 Vand. L. Rev. 1247 (2014); Rory Van Loo, *Stress Testing Governance*, 75 Vand. L. Rev. 553 (2023). Continuous adversarial testing describes, simply, adversarial testing that is iterative. This concept exists for ML training specifically. *See* Mohamed elShehaby & Aditya Kotha, *Introducing Adaptive Continuous Adversarial Training (ACAT) to Enhance ML Robustness*, arXiv (March, 15, 2024), https://arxiv.org/html/2403.10461v1.

Importantly, adversarial testing should not become a one-time hurdle or transform into a seal of approval that suggests all risks have been reviewed. Instead, the process should remain dynamic, with regular reassessment as technologies are deployed in new contexts and for new uses. An iterative approach allows regulators to discover emerging harms, address varying impacts across different populations, and identify new attack vectors that may not have been available during initial testing.

Effective oversight also cannot depend on voluntary arrangements alone. Mandatory requirements for early access and testing before deployment will enable regulators to implement necessary restrictions and friction points where and when they matter most, rather than attempting to modify systems that are already operational and widely deployed.

Human Subjects Oversight

The rapid and constant deployment of untested technologies into a range of services and onto unsuspecting audiences has effectively turned users into involuntary participants in corporate experiments. Being a test subject should not be a condition of human life. A suite of policy mechanisms optimized for oversight of disaggregated, data-driven processes–whether those processes involve primarily personal data or other kinds of data–can borrow from accumulated learning on how to structure oversight of human subjects testing. Today, oversight of government-funded human subjects research is anchored in the Belmont Principles, as follows: First and above all, do no harm (democratic legitimacy); second, to the greatest extent feasible, preserve research subject autonomy; third, ensure reciprocity and cost/benefit balance.[29] Those principles guide our recommendations here.

Adapting the existing human subjects toolkit to disaggregated, data-driven architectures and processes requires some translation, in two important and interrelated ways. First, consent plays (and should play) a significant role in human subjects research, but the human subjects toolkit also includes mechanisms for putting consent in its proper place. Sometimes, human subjects can consent to participate in research, but only when consent is feasible. Conversely, in some cases where obtaining consent would be impracticable or even directly counterproductive, it is permissible not to inform human test subjects or even to deceive them. In such cases, however, the human subjects toolkit inserts oversight at the center. During the design phase, both research aims and research methods must be screened and approved, and researchers must answer important questions about possible harms to both individuals and subject populations more generally.

Second, the human subjects toolkit is designed for oversight of discrete experiments, but, although operators of distributed digital processes sometimes construct discrete experiments, they also design and superintend processes of continuous, data-driven optimization. Those processes are tunable based on externally determined parameters, which, as things now stand, are set based on metrics attuned to profit and, in an ad hoc fashion, political pressure and personal beliefs of those in a position to control these processes. Mechanisms for human subjects oversight therefore must include requirements that relate specifically to optimization processes, parameters, and tuning.

29 *See* The Nat'l Comm'n for the Prot. of Human Subjects of Biomedical & Behav. Rsch., The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research (1979); 21 C.F.R. Ch. I, Part 56.

Drawing on the Belmont Principles while addressing contemporary challenges, we propose three core elements of an improved human subjects oversight framework for networked digital architectures, systems, and processes. The first element involves mandatory registry and disclosure of experimental protocols and processes. It would require documentation of all experimental and/or continuous optimization processes involving human subjects and their parameters, including details about the variables, how they are applied and to which populations, and any applicable experimental protocols and hypotheses. The registry would function as a crucial governance seam, enabling oversight and additional regulation focusing on the human impacts of networked digital processes and services. Following a tiered model similar to FDA clinical trial disclosure, the registry would provide different levels of access for regulators, researchers, and the public, creating essential points of visibility while protecting legitimate business interests.

The second element involves establishing regulatory authority over both experiments and optimization processes, including both the power to set limits and the power to ban certain practices entirely. Regulators would be able to impose limits on A/B testing and continuous machine-driven optimization, to require adjustments to optimization parameters based on emerging evidence about impacts, and/or to halt harmful experiments or optimization processes when necessary.

The third element introduces enhanced requirements for large-scale systems. Platforms, services, and models above certain size thresholds would face stricter oversight, including the mandatory use of control populations for all experiments and optimizations and enhanced monitoring and reporting requirements (all of which are comparable in some ways to the testing requirements imposed by the FDA before a new drug may be released to the general public or an existing drug may be approved for new uses[30]). This tiered approach recognizes that systems with broader reach require more careful scrutiny and more robust safeguards. Ideally, the friction generated by this regime would operate as a kind of tax that would land most heavily on forms of experimentation that are socially wasteful because they produce harmful impacts.

---

30 *See* Anna Lenhart & Sarah Myers West, *Lessons from the FDA for AI*, AI Now Inst., https://ainowinstitute.org/lessons-from-the-fda-for-ai; *see also* Andrew Tutt, *An FDA for Algorithms*, 69 Admin. L. Rev. 83 (2017) (proposing the creation of an agency modeled on the FDA to oversee algorithmic safety).

Part 4: Institutional Mechanisms for Implementing Regulatory Dynamism

The policymaking toolkit discussed above requires a corresponding institutional framework, meaning new statutory authorities to replace the relevant parts of the Administrative Procedure Act (APA) and their corresponding agency implementations. This presents a significant institutional design challenge. On one hand, regularized frameworks are necessary for both internal/political and public/democratic accountability; on the other, such frameworks can lend themselves to industry capture, political overreach, and/or process paralysis. And yet the particular procedural frameworks described in the APA and extended by agency framework statutes–which are deeply ingrained in the contemporary legal consciousness–are not the only ways that administrative processes might be designed.

As a general matter, we think the principles described in Part 2 can point the way toward the kinds of institutional redesign that are needed. Administrative processes for the networked information economy must permit rapid (but orderly) iteration of policymaking actions, but prioritizing speed for its own sake would be indefensible. Instead, regulatory "efficiency" must be judged relative to a complex set of values that includes accountability to the public and successful implementation of mandates designed to serve public goals and priorities. Administrative processes for the networked information economy therefore require their own governance seams to enable internal and public oversight, and they must be designed in ways that hold administrative efficiency and necessary process frictions in balance. Here, we begin to sketch important elements of such processes but leave other important questions for future modules of this project to address.

Streamlined, Iterative Rulemaking

One essential baseline for dynamic and effective policymaking involves a streamlined but iterative rulemaking process. As noted in Part 1, the formal rulemaking process defined in the APA has largely fallen into disuse; where versions of it survive–most notably for our purposes, in the special rulemaking requirements that apply when the FTC seeks to exercise its unfair and deceptive acts and practices authority proactively–it has effectively prevented agencies from engaging in innovative policymaking in the public interest.[31] We recommend that both the general provisions for formal rulemaking and specialized formal rulemaking processes such as those governing the FTC be eliminated. Provisions for informal rulemaking, meanwhile, should be revised with two goals in mind.

First, informal rulemaking (in our preferred world, just "rulemaking"), should be fortified more effectively against the kinds of capture strategies that produce process paralysis. Without question, agencies should listen to the public and should receive input from those whom their choices will affect (our forthcoming report on inclusion mechanisms addresses this requirement in detail). But agencies also should be equipped with both the resources (including information and analytical capacity) to receive public input effectively and the leeway to defer regulated industry objections until subsequent iterations of the rulemaking process have made the record clearer.[32]

31 Magnuson-Moss Warranty--Federal Trade Commission Improvement Act, Pub. L. No. 93-637, § 202, 88 Stat. 2183, 2193-95 (1975); Chris Jay Hoofnagle, Federal Trade Commission Privacy Law and Policy 56, 334 (2016).

32 *See* Steven J. Balla, et al., *Responding to Mass, Computer-Generated, and Malattributed Comments*, 74 Admin L. Rev. 95 (2022).

Second, rulemaking should be iterative by default. By this we mean: that rules should provide for periodic updates; that agencies should both be held responsible for updating issued rules and given the resources necessary to do so. We do not mean to suggest, however, that issued rules should sunset if not updated, and indeed we think such a requirement would worsen process paralysis. We mean, rather, exactly the opposite: that agencies should be empowered to make rules and then update them periodically without laboring uphill under the "regulatory burden" mindset described in Part 1 each time they act. In particular, many of the interventions we describe in Part 3 will require periodic updates, transparent criteria for modification of applicable requirements, and integration support encompassing technical specifications for implementation, testing frameworks for compliance, documentation requirements for different scales of operation, and resources for smaller actors in the ecosystem. To ensure that data flow restrictions, design requirements, adversarial testing processes, and human subjects restrictions remain effective, regulators will need protocols for regularly assessing impacts and implementation challenges across different timescales and contexts.

Interdisciplinary Framing and Assessment

A second kind of essential baseline for dynamic and effective policymaking involves regularized structures for interdisciplinary problem framing and assessment. The complex socio-technical challenges presented by networked digital architectures implicate multiple disciplines, including computer science, human-computer interaction, technology studies, sociology and economic sociology, behavioral economics, anthropology, and applied ethics. Perspectives from experts in these disciplines can assist regulators in determining where and how to create governance seams, designing and implementing feasible and appropriate friction mechanisms, and/or devising experimental protocols to test those mechanisms before mandating their ongoing implementation. Technical experts can advise on implementation feasibility; behavioral scientists can evaluate user impacts, identify unintended consequences, and predict future challenges; and social scientists and ethicists can identify and evaluate systemic implications. In similar fashion, perspectives from multiple disciplines can inform the operation of human subjects oversight panels and the development of design requirements and data flow restrictions. In continuous adversarial testing, interdisciplinary experts can help regulators study how patterns and behaviors drift through model training, deployment, and retraining.

Rather than relying on siloed pockets of narrow expertise or temporary consultations (that may also become additional sites of process paralysis), agencies should have the authority and the resources to hire in-house interdisciplinary teams and integrate them directly into policymaking processes from the outset. Those teams, in turn, can be empowered to develop assessment frameworks that explicitly incorporate multiple disciplinary perspectives and establish processes for resolving conflicts between different disciplinary approaches. Here again, Paperwork Reduction Act mandates that currently would operate as straitjackets on such processes should be revised and rethought to enable beneficial collaboration without reintroducing new types of information overload.

Here again, it is important to make clear what we do not mean to suggest. This recommendation is being offered at a moment when public anger with seemingly unaccountable expert authority is high. We believe that subject matter expertise has an important role to play in the project of devising appropriate oversight measures for the networked information economy, but we also believe that experts should serve, listen to, and collaborate with the public. Another module in this project will focus on the inclusion of affected publics in agency policy making and will offer recommendations for fostering more comprehensive and genuine communication and collaboration between the public and subject matter experts.

Experimental Protocol Design and Tuning

The additional, experimentalist processes described in Part 3 may sometimes require more specialized structures for defining goals and methods and evaluating performance. Such structures must accommodate the trial-and-error nature of experimental regulation by allowing for initial, experimental parameterization, intermediate review and assessment, and iterative tuning.

The universe of administrative procedure contains partial analogies for the sort of thing that is required. Many agencies, for example, can issue requests for information (RFIs) to elicit comments and suggestions regarding a contemplated rulemaking (and to help create a record in case the final rule is contested). Some agencies, such as the Consumer Financial Protection Bureau and the Federal Transit Authority, are empowered to create so-called regulatory sandboxes in which regulatory oversight is relaxed in some respects to permit regulated entities to experiment in certain ways that are thought to promise public benefits.[33] As these examples suggest, sandbox experiments generally have expressed a deregulatory orientation consistent with the understanding of regulation that we described in Part 1, above. But the sandbox mechanism also might be deployed in aid of heightened regulatory supervision.[34]

Still other structures within the existing administrative state are designed to enable experimentation, but in a way one degree removed from agency policymaking. So, for example, NIST has established the Assessing Risks and Impacts of AI (ARIA) program to support the evaluation, red teaming, and field testing of AI models and systems submitted by their developers.[35] The program will develop consensus methods, guidelines, and metrics that AI developers and other organizations can use to perform their own evaluations. Other agencies, in turn, can refer to the results of the ARIA program and other NIST processes in their policymaking activities (and, as we discuss in our concept paper on Provisioning Digital Tools and Systems for Government Use, they may be required to rely on NIST standards in their procurement and systems development activities).

Informed by these examples, Congress could authorize agencies to establish what we call policy sandboxes to develop, test, and tune experimental protocols for oversight of networked digital architectures, systems, and processes. An arrangement modeled in part on the RFI could invite comments from regulated entities and the public about the planned program of experimental oversight the agency intends to undertake. Procedural requirements that attach at this phase should focus on development and documentation of the experimental protocol following the initial authorization and could be informed by the processes in use at NIST. Documentation requirements should cover the high-level policy justifying creation of the policy sandbox, the technical parameters being tested and tuned, and the frequency and method(s) of iterative assessment. Last and importantly, the rigid mandates of the Paperwork Reduction Act should be revised to enable creation and iterative documentation of protocols for regulatory experimentation such as those described here.

Implementing policy sandboxes effectively would require expanded monitoring and enforcement capabilities. In particular, regulators need to be able to identify and respond to the kinds of workarounds that networked digital environments can enable. The proposals in our concept paper on regulatory monitoring are designed to create the kinds of visibility that would be required. A future module of this project will address enforcement and corrective measures.

---

33 *See* 12 C.F.R. § 1002.8 (2023); *Mobility on Demand (MOD) Sandbox Program*, U.S. Dep't of Transp., https://www.transit.dot.gov/research-innovation/mobility-demand-mod-sandbox-program.

34 *See generally* Hilary Allen, *Regulatory Sandboxes*, 87 Geo. Wash. L. Rev. 579 (2019); Douglas Sarro, *Sandbox Fictions*, __ Osgoode Hall L. Rev. (forthcoming 2025).

35 *ARIA Resources*, Nat'l Inst. Standards & Tech., https://ai-challenges.nist.gov/aria/library.

Because this recommendation is being offered at a moment in which federal agencies are being confronted with a different kind of experimentalism, it is important to make clear what we do not recommend. In our view, the requirement of prior legislative authorization is non-negotiable. Such authorization could occur on a case by case basis, as part of the framework statute for a new agency such as the oversight board described in our concept paper on regulatory monitoring, or as part of a substantial overhaul of the APA. In addition, the mechanism we envisioned would authorize only public regulatory experimentation—meaning that the authorization to develop and operate policy sandboxes should clearly prohibit agencies from using either the procurement process or the Federal Advisory Committee Act to devolve experimentalist authority to any private vendor or consultant or to a committee comprised largely of stakeholder representatives and other private actors.

Iterative Certification and/or Licensing

Regulators will need tools and processes for verifying that covered entities are complying with both the requirements established via iterative rulemaking and any additional requirements and tuning parameters imposed via a human subjects protocol or in a policy sandbox. In part, this is a familiar compliance problem, but addressing it requires new thinking about how to verify compliance without reproducing the dysfunctions described in Part 1. Our concept paper on regulatory monitoring addresses that question. Addressing compliance failures effectively also may require new enforcement mechanisms, a topic to be explored in a subsequent module of this project.

In some cases, premarket certification and/or licensing of particular systems or processes will be justified based on the threat of harm that uncontrolled development poses for public health, public safety, and/or the security and stability of critical infrastructures and systems. So, for example, the Food, Drug and Cosmetic Act requires premarket approval of both drugs and certain classes of medical devices. The process, superintended by the Food and Drug Administration, involves multiple steps, including an investigational new drug application notifying the FDA of intent to engage in human testing, a new drug application incorporating information learned from clinical trials and including various other categories of information (for example, manufacturing processes, product description, quality control and assurance procedures, proposed labeling, etc.), the possibility of conditional approval (including a requirement for post-approval studies), and the possibility of expediting the review process in appropriate cases.[36] The Federal Aviation Administration certifies the airworthiness of new aircraft designs and major changes to aircraft designs using a multi-step process that incorporates testing of major components, flight testing, and mandatory, prompt reporting of failures and malfunctions.[37] A complex regulatory framework governs all aspects of nuclear plant construction and operation.[38] The Nuclear Regulatory Commission certifies the initial reactor design and control systems, then requires comprehensive logging of any changes to the approved parameters, and finally continually verifies that the plant is still operating within its licensed specifications.

By analogy, Congress could create or authorize regulators to develop new kinds of licensing regimes adapted to the particular circumstances of digital architectures, systems, and processes.[39] Different requirements might be developed for different circumstances. For example, new applications of existing AI models might be subject to one type of licensing, social media recommender systems and real-time ad auction systems to others, and large language models and certain other automated decision-making models to yet others.

---

36 For detailed consideration of the FDA as a possible template for AI regulation, see Anna Lenhart & Sarah Myers West, *Lessons from the FDA for AI*, AI Now Inst., https://ainowinstitute.org/lessons-from-the-fda-for-ai.

37 4 C.F.R., Ch. I, Part 21.

38 10 C.F.R. § 50; *Backgrounder on Nuclear Power Plant Licensing Process*, U.S. NRC, https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/licensing-process-fs.html.

39 For a similar proposal for the creation of a new agency to implement a licensing regime for complex financial instruments, see Saule T. Omarova, *License to Deal: Mandatory Approval of Complex Financial Products*, 90 Wash. U. L. Rev. 63, 66, 129 (2012).

Once again, there are caveats. Both licensing and continuing certification processes can be very slow. As a general matter, it will be important to design processes capable of operating more nimbly; however, in order to avoid the problems of process paralysis and checkbox compliance described in Part 1, above, it will also be important to design processes that operate with rigor. We think part of the answer here involves equipping agencies with the technical and financial resources necessary to engage in real-time monitoring and assessment of digital architectures, processes, and systems, both to assess initial eligibility for licensing and to verify continuing eligibility. Regulated entities also should be obligated to self-monitor and report newly discovered risks and incidents. Part of the answer also likely involves a restructured menu of sanctions for noncompliance. Additionally, it is worth underscoring that both unrestricted private development and public licensing are inherently imperfect activities. When public health, public safety, and the integrity of critical systems are at stake—or when the risks created by a networked system or process are systemic, emergent, and potentially very high—it is worth proceeding with some caution.

Situating Oversight and Contestation within the Experimentalist Framework

Regularized structures for experimental regulation also must enable regulators to be held accountable both internally and to the public. The two kinds of accountability, moreover, are not interchangeable. Internal control without public accountability can enable political overreach. At the same time (as our govtech report also illustrates), accountability mandates can become straitjackets, so the question is how to enable oversight and contestation without bringing ostensibly nimble, experimentalist processes grinding to a halt.

Here, some of the most relevant analogies sit outside the framework established by the APA and reiterated, with some variation, in the statutes creating and empowering the various executive branch and independent agencies. As already noted, in industries that follow iterative, experimental protocols, those processes are internally accountable in various ways. Staff charged with experimentation still must satisfy requirements of documentation and accountability to managers. And, in certain regulated industries, ongoing dialogues with regulators about the need to satisfy public mandates have given rise to rich sets of conventions for documenting experimental work.[40] Similarly, in both basic science and clinical research, requirements of reproducibility and falsifiability demand documentation and publication of the methods used.[41] Lessons learned across these varied domains could inform the development of a framework statute establishing documentation and review requirements for experimental regulation.[42]

As experimental activities have encountered public regulatory mandates, hybrid mechanisms for overseeing the design and conduct of experimental research also have emerged. In particular, human subjects review boards require description and, as necessary, justification of experimental methods and protocols. And the constitu-

40 *See, e.g.,* U.S. Food & Drug Admin., *Data Integrity and Compliance with Drug CGMP: Questions and Answers: Guidance for Industry*, 83 Fed. Reg. 64,132 (Dec. 13, 2018); OECD, *Position Paper on Good Laboratory Practice and IT Security*, ENV/CBC/MONO(2021)26 (Sept, 2021), https://www.oecd.org/en/publications/oecd-position-paper-on-good-laboratory-practice-and-it-security_910b7bd2-en.html; Huan Gui, Ya Xu, Anmol Bhasin & Jiawei Han, *Network A/B Testing: From Sampling to Estimation, in* Proc. 24th Int'l Conf. World Wide Web 399 (2015); Jinfang Sheng, Huadan Liu, & Bin Wang, *Research on the Optimization of A/B Testing System Based on Dynamic Strategy Distribution,* 11 Processes 912 (2023); Krishan Kumar., *Good Documentation Practices (GDPs) in Pharmaceutical Industry,* 4 J. Anal. Pharm. Res. 00100 (2017); Tim Sandle, *Good Documentation Practices,* 20 J. Valid. Technol. 1 (2014); Test Procedures, Nat'l Highway & Traffic Safety Admin., https://www.nhtsa.gov/vehicle-manufacturers/test-procedures.

41 *See, e.g.*, C. Glenn Begley & John P.A. Ioannidis, *Reproducibility in Science: Improving the Standard for Basic and Preclinical Research*, 116 Circ. Res. 116 (2015).

42 As an example of this type of work, see National Institutes of Health, *Guidelines and Policies for the Conduct of Research in the Intramural Research Program at NIH* (8th ed. 2023), https://oir.nih.gov/sourcebook/ethical-conduct/research-ethics/nih-guidelines.

tion and ongoing operation of human subjects review boards is itself subject to various criteria of adequacy.[43] Some participants in regulated industries also have developed criteria for assessing the efficacy of particular research programs, based, for example, on criteria such as patent outcomes, health outcomes, or consumer feedback.[44] In social science, researchers have begun to develop criteria for assessing research impact.[45] Lessons learned across these varied domains could inform the development of a framework statute establishing requirements for the design and ongoing assessment of programs of experimental regulation.

Last and importantly, a mandate empowering regulators to conduct experimental regulation should provide for both appropriate internal oversight and external challenge in the courts. These proposed requirements intersect with larger separation of powers and due process questions and therefore sit mostly outside the scope of this concept paper, so we mention them only briefly here. First, documentation and process requirements should enable agency inspectors general to reconstruct experimentalist processes after the fact, evaluate their sufficiency, and/or direct changes to experimental protocols in process. For their part, inspectors general should be charged to distinguish between incidental and/or unintentional deviations from best practices and systematic flouting. Second, documentation and process requirements for experimental regulation need to be designed in ways that create public-facing visibility and permit contestation in the courts. All of these requirements, however, need to be crafted in a way that is attentive to the twinned risks of process paralysis and checkbox compliance. A subsequent module of this project will address the rule-of-law requirements that should attach to redesigned, information-era administrative processes.

Resources for Regulatory Dynamism

In our two previous concept papers, we have observed that comprehensive redesign of administrative processes requires equipping the administrative state with the necessary resources, including especially personnel with appropriate skills. We repeat those recommendations here.

We also want to underscore a more fundamental mismatch between our recommendations for new mechanisms and processes and prevailing budgeting practices, which (when they function at all) envision an annual funding cycle. Like the recommendations in our two previous reports, many of those developed here will require investments that span annual funding cycles. Changes to the budget process designed to facilitate long term budgeting are outside the scope of this paper, but we agree with those who have argued that the issue requires sustained attention.[46]

---

43 *See* 21 C.F.R. Ch. I, Part 56; Lura Abbott & Christine Grady, *A Systematic Review of the Empirical Literature Evaluating IRBs: What We Know and What We Still Need to Learn*, 6 J. Empirical Res. on Hum. Res. Ethics 3 (2011).

44 *See, e.g.,* Annmarie Kelleher, et al., *Consumer Feedback to Steer the Future of Assistive Technology Research and Development: A Pilot Study*, 23 Top Spinal Cord Inj. Rehab. 89 (2017); Mark A. Lemley and Carl Shapiro, *Probabilistic Patents*, 19 J. Econ. Persp. 75 (2005); Elizabeth Gargon, et al., *Choosing Important Health Outcomes for Comparative Effectiveness Research: A Systematic Review*, 9 PLoS One e99111 (2014); Sara Ahmed, et al., *The Use of Patient-Reported Outcomes (PRO) within Comparative Effectiveness Research: Implications for Clinical Practice and Health Care Policy*, 50 Med. Care 1060 (2012).

45 For a useful overview, see Saba Hinrichs-Krapels & Jonathan Grant, *Exploring the Effectiveness, Efficiency and Equity (3E's) of Research and Research Impact Assessment*, 2 Palgrave Comm'c'ns 16090 (2016).

46 For useful discussions, see Alice M. Rivlin, *Rescuing the Budget Process*, 32 Pub. Budgeting & Fin. 53 (2012); Megan Lynch, Cong. Rsch. Serv., R45789, Long-Term Budgeting Within the Congressional Budget Process: In Brief (2019).

Bibliography

Failures and critiques of the regulatory state

Julie E. Cohen, Between Truth and Power: The Legal Constructions of Informational Capitalism (2019).

Sidney Shapiro, Elizabeth Fisher & Wendy Wagner, *The Enlightenment of Administrative Law: Looking Inside the Agency for Legitimacy*, 47 Wake Forest L. Rev. 463 (2012).

Stuart Shapiro, *The Paperwork Reduction Act: Benefits, Costs and Directions for Reform*, 30 Gov't Info. Q. 204 (2013).

Jodi L. Short, *Regulatory Managerialism as Gaslighting Government*, 86 L. & Contemp. Probs. 1 (2023).

Regulatory innovation and experimentation

Hilary J. Allen, *Regulatory Innovation and Permission to Fail: The Case of Suptech*, 19 N.Y.U. J. L. & Bus. 237 (2023).

Cary Coglianese, *The Limits of Performance-Based Regulation*, 50 U. Mich. J. L. Reform 525 (2017).

Brett Frischmann & Paul Ohm, *Governance Seams*, 37 Harv. L. J. & Tech. 1117 (2023).

Brett Frischmann & Susan Benesch, *Friction-in-design Regulation as 21st Century Time, Place, and Manner Restriction*, 25 Yale. L. J. & Tech. 376 (2023).

Ellen P. Goodman, *Digital Fidelity and Friction*, 21 Nev. L.J. 623 (2021).

Paul Ohm & Jonathan Frankle, *Desirable Inefficiency*, 70 Fla. L. Rev. 777 (2018).

Charles F. Sabel & William H. Simon, *Minimalism and Experimentalism in the Administrative State*, 100 Geo. L. J. 53 (2011).

Lauren Willis, *Performance-Based Consumer and Investor Protection: Corporate Responsibility Without Blame, in* The Culpable Corporate Mind (Elise Bant, ed., 2022).

Governance challenges of digital architectures, processes, and systems

Kenneth A. Bamberger, *Technologies of Compliance: Risk and Regulation in a Digital Age*, 88 Tex. L. Rev. 660 (2010).

Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, 3 Big Data & Society 1 (2016).

Jennifer Cobbe, Michael Veale, Jatinder Singh, *Understanding Accountability in Algorithmic Supply Chains, in* Proc. ACM FAccT Conf. (2023).

Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a 'Right to an Explanation' Is Probably Not the Remedy You Are Looking For*, 16 Duke L. & Tech. Rev. 18 (2017).

Alex Moehring, et al., *Better Feeds: Algorithms That Put People First*, Knight Georgetown Inst. (Mar. 2025), https://kgi.georgetown.edu/research-and-commentary/better-feeds/.

Ranjit Singh, *Good Technology Is Slow (to Scale), in* The Good Robot: Why Technology Needs Feminism 100 (Eleanor Drage & Kerry McInerney, eds., 2024).

Building the dynamic regulatory toolkit

Lura Abbott & Christine Grady, *A Systematic Review of the Empirical Literature Evaluating IRBs: What We Know and What We Still Need to Learn*, 6 J. Empirical Res. on Hum. Res. Ethics 3 (2011).

Hilary J. Allen, *Regulatory Sandboxes*, 87 Geo. Wash. L. Rev. 579 (2019).

Mehrsa Baradaran, *Regulation by Hypothetical*, 67 Vand. L. Rev. 1247 (2014).

Ayelet Gordon-Tapiero, Paul Ohm & Ashwin Ramaswami, *Fact and Friction: A Case Study in the Fight Against False News*, 57 UC Davis L. Rev. 171 (2023).

James Grimmelmann, *The Law and Ethics of Experiments on Social Media Users*, 13 Colo. Tech. L.J. 219 (2015).

Johanna Gunawan, Woodrow Hartzog & Neil Richards, *Dark Pattern as Disloyal Design*, 99 Ind. L. J. (forthcoming 2025).

Woodrow Hartzog, Privacy's Blueprint: The Battle to Control the Design of New Technologies (2018).

Anna Lenhart & Sarah Myers West, *Lessons from the FDA for AI*, AI Now Inst. (Aug. 1, 2023), https://ainowinstitute.org/lessons-from-the-fda-for-ai.

Saule T. Omarova, *License to Deal: Mandatory Approval of Complex Financial Products*, 90 Wash. U. L. Rev. 63 (2012).

Rory Van Loo, *Stress Testing Governance*, 75 Vand. L. Rev. 553 (2023).

ITT REDESIGNING THE GOVERNANCE STACK

GEORGETOWN LAW