

Building On Regulatory Foundations

& Bridging To The Future



Summaries of Breakout Discussions

- The Role of Technology in Public Participation in the Regulatory Process
- Is OIRA Still Fit for Purpose?
- Applications of Generative AI to Rule Development and Evaluation
- Improving Evidence for Prospective & Retrospective Policy Analysis
- Agile Regulation

Introduction

Modern U.S. regulatory practices have their roots in the first regulatory agencies established in the late 19th century, in the Administrative Procedure Act of 1946, and in the regulatory reforms of the 1970s and 1980s.

Two more recent, but no less important, foundations are Executive Order 12866 of 1993 and Circular A-4 of 2003. In November 2023, the **IBM Center for the Business of Government** and the **George Washington University Regulatory Studies Center** gathered regulatory experts to reflect on these foundations and explore emerging opportunities and challenges for good regulatory practice.

Daniel Chenok, director of the IBM Center for The Business of Government, and Susan Dudley, founder of the GW Regulatory Studies Center welcomed the 100+ participants and introduced keynote speaker, Cass Sunstein of Harvard Law School. Sunstein focused his remarks on what he calls “sludge,” the often-unnecessary red tape that can reduce efficiency and limit access to government services.

Following these introductory remarks, participants divided into five breakout sessions for more in-depth discussions of the following topics. The sessions were held under the Chatham House Rule, where participants are free to use the information received, but are not allowed to reveal who made any particular comment. Learn more about the event at <https://regulatorystudies.columbian.gwu.edu/event-highlights-building-regulatory-foundations-and-bridging-future>.

The Role of Technology in Public Participation in the Regulatory Process

This breakout discussion explored the role of technology—particularly artificial intelligence (AI)—in public participation in the regulatory process. Steve Balla, professor of political science at the George Washington University, facilitated the topic and participants considered all types of public participation. The conversation often circled back to public commenting as a key form of participation in the rulemaking process. In that domain, participants identified ways in which AI could support and challenge current practices.

Furthermore, participants shared that AI could help individuals shape their public comments into a format most useful to agencies. For example, AI could transform a verbal interview with a regulated party into a written comment, accounting for what agencies want to see from public comments. This sort of tool would let individuals speak in their natural language and convey their points as they know how, and the AI could help reshape their contributions into the format agencies need. This could lift some of the burden on individuals and help their comments be more useful for agencies.

AI could also potentially support an online tool that reviews individuals' written comments and offers suggestions for transforming preference-based comments to substantive comments. Statements of sentiment or preference are not as useful to agencies, which must base their regulations on the substantive administrative record. A tool to push users to justify their preference by providing their unique knowledge and lived experience could make a difference in how agencies respond to their comment, and whether their perspective affects the final rule.

While AI could certainly improve public commenting if used to support tools such as these, it could also introduce new challenges. For one, would AI dilute substantive comments on big rules? Generative AI—like ChatGPT—gives users a platform to create an individualized public comment with minimal effort. Individuals could use generative AI to quickly write and submit multiple unique comments on a given proposal. While agencies have technology to identify duplicate comments, their technology may not be able to identify and filter out similar but unique comments. More unique comments might make it more challenging for agencies to address all the comments they receive and enact rules.

The group also discussed considerations beyond whether AI is helpful or harmful for public commenting. Attendees agreed that different tools may be relevant depending on the nature of the comment and characteristics of the proposed rule. With respect to comments: some are substantive—providing data, economic information, and situated knowledge—while others are preference-based, only offering whether they agree or disagree with the proposal. Regarding the nature of the rule: does the rule receive many comments, or only a handful of comments?

Because of how diverse both comments and rules can be, it may be appropriate to have a variety of technological tools to best assist in the different circumstances.

Building On Regulatory Foundations - Engagement

Session attendees also discussed whether a lack of technology is a barrier to public participation, at all. Rulemaking is an inherently legal process. One session attendee hypothesized that the legal nature, the lengthy and complicated process, and a lack of access to the process are the primary barriers to effective public participation, rather than a lack of technological tools. An average American may not be able to engage as effectively with the process as a sophisticated commenter, like a corporation or an advocacy group, that has repeat experience with the regulatory process and resources to devote to writing effective comments. Given this discussion, the group reflected on whether public commenting is the most effective strategy for public participation in the regulatory process. By the time proposed rules are available for comment, agencies have invested significant time and resources into its development and may be hesitant to make significant changes.

Broadly, the goal of public engagement in the regulatory process is to create rules that are responsive to people's needs. How do policymakers create an interactive, responsive system of public participation? This panel largely agreed that public commenting, as it currently exists, is not sufficient for effective public participation in rulemaking. Whether technologies like AI will help to create that interactive, responsive system remains to be seen.

Is OIRA Still Fit for Purpose?

A second breakout discussion focused on whether the Office of Information and Regulatory Affairs (OIRA), created in 1981, is still fit for purpose. Bridget Dooling, currently an assistant professor of law at the Ohio State University and previously a research professor with the George Washington University Regulatory Studies Center and deputy OIRA branch chief, led the discussion.

Attendees of this breakout session debated a series of interrelated questions: What is OIRA's purpose? Has its purpose changed over time? And is OIRA currently capable of fulfilling its purpose?

There was general consensus regarding OIRA's original purpose: Its primary function was to serve as a rational bureaucracy, staffed by apolitical generalists capable of mediating between competing interest groups within the government in order to produce regulations that yield the greatest net social benefit. OIRA desk officers must evaluate the analyses underlying agencies' regulatory proposals against fundamental economic and statistical principles. OIRA's secondary

function was to facilitate coordination and oversight in the regulatory process, so as to provide greater transparency and accountability to the elected president.

The breakout session attendees also agreed that OIRA's purpose has changed over time in two major ways. The first significant development is a consequence of the increasingly technical content of the proposed rules OIRA is tasked with reviewing. In response to this complexity, OIRA has been forced to adopt a more specialized function than it was originally intended to fulfill. The office can no longer assess proposed regulations by relying solely on economic and statistical analyses. Rather, it must now supplement those core competencies with deep analyses of scientific and technological subjects.

The second change involves the politicization of the office. While OIRA—created by President Reagan—was initially perceived as having a deregulatory agenda, the perpetuation of the office in the Clinton administration gave it a nonpartisan reputation. The apolitical nature of the OIRA staff enabled them to provide regulatory continuity over the years, despite the turnover of administrations. Breakout attendees worried that OIRA has become more politicized in recent years, harming its image and detracting from its primary analytic function. While there will always be a tension between OIRA's role as an analytical reviewer and its role as a representative of the elected president, participants thought it could more effectively fulfill its purpose if it were perceived as less political.

Breakout participants were split evenly on the question of whether OIRA currently has the capacity to fulfill its purpose. Those who believed that OIRA is still fit for purpose highlighted its stellar analytic capabilities, its wealth of institutional knowledge, the time-tested role that it performs within the executive branch, the esteem in which the president holds the office, and its archetypal position within the international regulatory community.

Other attendees believed that OIRA could improve in various ways. Some thought that in comparison to the regulatory agencies of some other countries, OIRA's approach to regulation insufficiently fosters technological progress. More specifically, they highlighted the pro-innovation regulatory approach used in South Korea which involves innovation reviews of entire fields, such as autonomous vehicles. Participants noted that OIRA would require additional staff to develop similar capabilities. A related but distinct critique of OIRA was a perceived lack of an aspirational global regulatory vision on topics of universal importance (e.g. climate change) and its corresponding failure to network adequately with the regulatory agencies of other countries.

While attendees lauded the capabilities of the current OIRA staff, others worried that the office is still too small to handle the deluge of draft rules sent to it for review. As of November 2023, OIRA has approximately 67 full-time employees, a number which has remained largely

unchanged over the years as the demands placed upon the office have dramatically increased. Concerned attendees proposed that OIRA would benefit from the addition of more subject matter specialists, particularly in burgeoning fields such as artificial intelligence. Others voiced their concern that increasing OIRA's staff too much could diminish its nimbleness and responsiveness.

Some argued that OIRA does not proactively communicate with other federal agencies and academic institutions. OIRA is viewed by many as an inscrutable black box, and consequently as an impediment to effective regulation. For this reason, several attendees recommended proactively teaching agencies and universities how OIRA functions and how they can constructively interact with it. This might entail training programs for members of other agencies and educational initiatives in both law and public affairs schools. Participants agreed that OIRA currently lacks an external constituency supporting its mission, and thought greater outreach could address that issue.

By the end of the breakout session, the attendees generally agreed with the following characterization of OIRA's trajectory:

The demands placed upon OIRA have changed throughout its lifetime. Nevertheless, the office has always managed to achieve its core analytic purpose, frequently adapting its approach in response to novel technical and political challenges. Whether OIRA can continue to fulfill its purpose in the contemporary regulatory landscape without increasing its staff and expanding the scope of its activity remains to be seen.

[Applications of Generative AI to Rule Development and Evaluation](#)

This discussion covered a variety of topics and several recurring themes emerged relating to the potential benefits and risks associated with the use of AI in the rulemaking process.

Dr. David Bray, Distinguished Fellow and Loomis Council Co-Chair at the Stimson Center, and Andy Fois, Chair of the Administrative Conference of the United States (ACUS), facilitated the generative AI breakout session, which was attended by participants from government agencies, academia, private sector companies, and non-profits. In general, the discussion primarily focused on AI's use during the public commenting stage of rulemaking, rather than another phase like proposal development or retrospective review. Specifically, much of the conversation centered on how agencies should respond to receiving comments that are entirely or partly generated by AI and how they might process comments using AI tools.

During a wide-ranging discussion, several prominent themes emerged. First, agencies use different approaches when addressing AI-generated comments, with consequential results.

The Federal Communication Commission’s (FCC) [Electronic Comment Filing System](#) and the General Services Administration’s (GSA) [eRulemaking program](#) (which is used by dozens of federal agencies as a shared service) represent the two broad strokes of public commenting processes. Participants discussed how back in the mid-2010s different government agencies experienced spikes in what appeared to be a mixture of human- and bot-submitted public comments, although an exceedingly small number of rulemaking efforts saw more than 10,000 comments. A 2021 analysis of a spike in public comments impacting the Environmental Protection Agency [observed](#) that “the 2002 E-Government Act did not anticipate the emergence of bots and thus fails to provide agencies with sufficient guidance on how to identify and treat bots and fake comments.”

Participants observed that the FCC had legally interpreted the Administrative Procedure Act (APA) of 1946 and related policies in a manner that effectively gave senior management less discretion to address the risk of comment surges – and precluded the FCC from being able to adopt the GSA’s eRulemaking program. During the mid-2010s, the then-CIO had attempted to make the case for the FCC to adopt the eRulemaking program and not succeeded. This stemmed from the FCC’s interpretations of the APA that prioritized real-time viewing of comments rather than waiting to post submissions until they are processed, acceptance of all comments even if they were perceived as potential spam, allowance of anonymous comments or comments with no identification checks, a reluctance to use CAPTCHA, and a strong push by external parties for the ability to submit comments in bulk. While the FCC’s legacy Electronic Comment Filing System eventually moved to a [cloud-based service](#) that included API rate limits for comment submissions, it employed a GSA service that before 2017 did not monitor API key requests for multiple registrations. After 2017, GSA’s public-facing platform, located at [regulations.gov](#), successfully implemented techniques such as CAPTCHA and API rate limits to [mitigate](#) the risk of being overwhelmed by automated submissions. The FCC since 2017 has made some adjustments too.

Participants also discussed how the 2017 net neutrality rulemaking demonstrated that the FCC’s interpretations of the APA made it technically at risk of [astroturfing](#) – defined as [organized activity](#) that falsely attempts to pass itself off as a grassroots movement. The Commission’s proposal received nearly 23 million comments in 2017, requiring the FCC to scale its cloud-based systems more than 3,000 percent to address the flood of comments. In 2021, the New York Attorney General identified at least [18 million](#) of these comments as not authentic. Since 2017, [regulations.gov](#) has not experienced the same issues, although [some rulemakings](#) routinely receive large volumes of mass submissions (though none that approach the scale of 23 million comments). Comparing different legal interpretations of the APA and the downstream impact on technical implementations highlighted how policy decisions may prevent the commenting process from being overwhelmed by bots or generative AI submissions.

A second theme was that generative AI presents opportunities for how agencies use it for processing comments. Participants mentioned the need for internal controls on processing and adjudicating comments, wondering how an agency should respond if a commenter claims their comment was not adequately considered by an AI system. One point that was brought up is that considering a comment is not akin to a vote or directly following its advice. For instance, an agency may consider the claims presented by critical comments and still determine that the proposed rule is worth pursuing regardless – public commenting does not overrule agency decisions. Nevertheless, the conversation routinely returned to the difficulty of defining what it means to both consider comments when AI is used and prove that the comments have been considered fully. In this instance, the FCC in the mid-2010s had a small success story in making all comments received during the 2014 and 2017 net neutrality discussions publicly available as a downloadable data set for others to analyze. As agencies work through this question, ensuring their policies for using AI comply with the APA and offering guidance and direction in interpretations and implementations is essential.

One participant suggested that agencies using AI systems decouple the procedures for processing comments from those for reviewing their substance to mitigate the issue of whether they have adequately considered public comments. While AI may serve a key role in categorizing, filtering, and summarizing public submissions, human involvement remains necessary in the latter stage of evaluating substantive feedback. On this topic, participants also deliberated about the role of transparency in how agencies adopt and use AI technologies. Like other technologies, relying on AI too much could distort the rulemaking process, so establishing boundaries and usage policies would aid public accountability. Generally, participants agreed that clear guidance from a body like the Office of Management and Budget on this matter would be valuable.

Third, participants considered on how to counteract potential negative effects of AI on the rulemaking process, while capturing its benefits for regulatory development. To mitigate risks, participants highlighted the value of red-teaming in predicting how bad actors might use AI technologies in the public commenting process. In fact, President Biden’s executive order from October 2023 on the development and use of AI incorporated “AI red-teaming” – [defined](#) as “a structured testing effort to find flaws and vulnerabilities in an AI system” – into several directives. Session participants mentioned how having a consistent process for handling public comments across agencies, built upon compatible interpretations of the APA, would make red-teaming activities even more valuable, since the results would apply in more contexts and fixing a vulnerability for one agency would fix it for all.

When discussing AI’s potential benefits to rule development, several participants posed whether such tools could help agencies formulate regulation. At the very least, multiple participants thought that AI could play a substantial role in synthesizing the data and science that inform

rules. Finally, individuals discussed how agencies can take proactive steps to better position themselves to reap the benefits of advanced AI capabilities. One example provided was how DOT's experience with using a [structured format](#) for its rules aided in retrospective review by facilitating the use of rule text as data. Greater accessibility of machine-readable rule text could play a role in [leveraging](#) AI systems to conduct ex ante and ex post analyses of regulations.

Ultimately, Generative AI poses numerous considerations for rule development and evaluation, particularly for public commenting on agency rules. Themes discussed in this session highlighted the need for agencies to be attentive and responsive to modern technologies and adapt their practices accordingly.

Improving Evidence for Prospective & Retrospective Policy Analysis

This breakout discussion focused on retrospective review of regulation, led by Nick Hart, President and CEO of the Data Foundation.

The session highlighted the important role of retrospective review in the regulatory process and acknowledged that more improvement is needed in this field. Attendees in the session dived into the risks and challenges for conducting retrospective review and discussed potential remedies for those challenges.

Dr. Hart pointed out that retrospective review is an essential part in the feedback loop of the rulemaking process. Retrospective review evaluates the effects of regulations after they have been implemented and verifies whether the existing regulations need to be removed, amended, or streamlined. Retrospective review also helps improve the design of future regulations. By comparing retrospective review against prospective regulatory impact analysis (RIA), regulators can assess whether the assumptions and models used in the ex-ante analysis are valid or need revisions. Such assessment will then inform the analysis and design of future regulations.

Despite the value of retrospective review, the consensus among session attendees was that agencies have not been doing a good job in retrospective review. There are at least [two major challenges](#). First, agencies lack incentives to perform retrospective review. Agencies are required to do ex-ante RIAs before they can issue new regulations, and the Office of Information and Regulatory Affairs (OIRA) serves as the gatekeeper to ensure that agencies comply with this requirement in the rulemaking process. After the regulation has been issued, however, the absence of retrospective review is less problematic for regulators. Regulators have little motivation to identify flaws in existing regulations or to disclose such information to the public. Moreover, there are often limited resources (i.e., capacity and time) within agencies, making regulators even less incentivized to evaluate their regulations retrospectively.

The second challenge is technical difficulties in designing the policy for evaluation. Once the regulation has been implemented, it is difficult to know what the counterfactual world (the one without the regulatory change) would have looked like. The design of regulation rarely takes into account data collection for retrospective review, so regulators often do not have the data needed for an ex-post analysis even if they intend to do so.

While agencies do not conduct retrospective review systematically, participants in the session discussed some effective practices. For example, the Environmental Protection Agency used regional rules as quasi-experiments to gather data for evaluating its gasoline standards. When it rescinded its Covid-19 vaccination mandate, the Occupational Safety and Health Administration was able to compare vaccine uptake after the rule was issued against the assumption made in the RIA of the rule. The EPA also has well-established data collection infrastructure for monitoring pesticides and incidents that may affect regulatory actions.

The session then explored potential remedies to achieve better systematic retrospective review. Some attendees suggested that OIRA can play a key role to provide leadership and guidance for retrospective review. Some mechanisms, such as automatic sunset of regulations, could also incentivize retrospective review. Under a sunset provision, agencies are required to review their rules by certain deadlines and, if they fail to do so, the rules will automatically expire.

Participants discussed the value of supporting internal or external reviews. For example, well-established economists supporting the development of prospective analysis and regulatory impact assessments may be [incentivized](#) to support retrospective review in ways program offices are not. Similarly, attendees pointed out that it does not have to be agencies who perform retrospective review. Such analysis can be contracted out to third-party analysts who may produce more objective evaluations.

It is also important to build data collection into the design of regulation such that data can be collected from the outset. One participant pointed to OIRA's new guidance on pilot projects and data collection in the revised [Circular A-4](#) issued in November 2023. The guidance recommends that agencies use pilot projects to test regulatory alternatives, if timing and other circumstances allow, and consider the regulatory alternatives that would facilitate data collection to support retrospective review.

The session concluded with a discussion about opportunities for artificial intelligence (AI) in retrospective review of regulation. For example, AI may alleviate the problem of resource limitations by helping regulators manage their work and time more efficiently. AI may also assist in collecting or generating the data necessary for analysis or mining existing data that are otherwise impossible to collect. While the lack of incentives persists, the progress in cutting-edge technologies will provide regulators with more tools to perform better retrospective review.

Agile Regulation

Agile regulation is a framework regulators can use to manage change at speed and scale. Michael Fitzpatrick, a partner at Brunswick Group, Senior Fellow at ACUS and Fellow at the National Academy of Public Administration (as well as a former OIRA deputy administrator), facilitated the session, which was attended by participants from government agencies, academia, private sector companies, and non-profits. In a world of rapid change, policymakers may have little time to adapt but are expected by market participants to respond quickly. While still a relatively new term, policymakers (including the [OECD](#)) increasingly refer to an agile regulation framework, a set of practices and strategies aimed at making regulators more flexible. This could be from agencies using agile processes in their own workflows or from building agility into regulations (e.g., [regulatory sandboxes](#)).

The session used the National Academy of Public Administration's paper [Agile Regulation: A Gateway to the Future](#), authored by a committee including Fitzpatrick, to launch discussion, with participants offering thoughts to critique or extend on points from the research. While participants discussed a range of potential agency actions, types of improvements roughly fell into three categories: internal processes, substantive design, and continuous learning. Some participants suggested that agency leaders could foster agile regulation by rewarding innovation among their employees. By highlighting the importance of adaptability and creative thought within trainings and incentivizing innovation in performance evaluations, leaders can signal the importance of these priorities to staff. Other participants noted that small teams often can act more nimbly than larger workforces and wondered if small strike forces focused on fast-moving problems in society could come up with efficient ways to respond. Multi-agency consortia could also share best practices to structure workforces and workstyles in a manner that encourages flexibility without ceding rigor (as well as shared information on newly emerging technologies). While not a major focus of the discussion, the group also discussed the potential of artificial intelligence to enhance agencies' internal processes. Among other tasks, AI could interpret data, generate or review agency ideas, or survey the corpus of past regulations for retrospective regulation to save staff time.

At the same time, some participants noted that agencies could not be expected to take on more responsibilities or oversight without appropriate resources. To reinvent or reform outdated processes, agencies might need financial support to hire external expertise, and integrating AI or retrospective review into workflows is likely to be costly.

Substantive design refers to how provisions within regulations are written, which can influence the actions regulated entities can take. Improving substantive design could increase agility by enabling market participants to experiment, lower costs, or move with greater speed. This

concept, however, led to discussions over whether and how agencies could change how they write rules without ceding authority to regulated entities. Some participants suggested that the relationship between regulator and regulated entity could be rethought, emphasizing partnership to a greater extent. Under this paradigm, agencies might be quicker to use rulemaking processes other than prescriptive or performance-based regulation. In some spheres, agencies have used innovations in substantive design to increase flexibility. Some pointed to the use of [adaptive licensing](#) by the Food and Drug Administration and thought lessons from that approval process could be applied in other areas. The [variety](#) of regulations surrounding autonomous vehicle testing shows that governments can build flexibility into regulation to advance beneficial technology.

Finally, by considering additional avenues for continuous learning, agencies may correct problematic rulemaking processes for the future and course correct past regulations that have had unexpected impacts. Some participants suggested actions that could enhance agencies' abilities to review their own effectiveness, particularly surrounding retrospective review. Multi-agency retrospective review could result in cost savings by spreading lessons learned across the administrative state, as well as leading to more honest evaluations because a rule examined by staff outside the author's agency could lead to more honest appraisals. Other participants wondered whether continuous commenting on issued regulation is feasible. If stakeholders could comment on regulations even after they are in effect, they could highlight ongoing problems and bring them to policymakers' attention. If agencies openly solicited feedback on operational rules, they could effectively enable continual retrospective review.

Participants had a spirited conversation, with several suggesting that the idea of agile regulation largely reiterated a need for innovation that many in agencies already recognized. However, most participants recognized that agencies may need to operate in new ways when faced with events that seem to be moving faster than ever, and all recognized the benefit of continual experimentation and evaluation among regulatory agencies.