REDESIGNING THE GOVERNANCE STACK
GEORGETOWN LAW

INSTITUTE FOR TECHNOLOGY LAW & POLICY
GEORGETOWN LAW

Regulatory Monitoring in the Information Economy

Preliminary Concept Paper

Julie E. Cohen

Brenda Dvoskin

Meg Leta Jones

Paul Ohm

Smitha Krishna Prasad

Co-signers:[1]

Woodrow Hartzog

Christopher Morten

Rory Van Loo

The harms of today's information economy–including deception and misinformation, pervasive state and commercial surveillance, and growing entrenchment and normalization of data-driven bias and inequality–are increasingly pervasive. Our administrative state, which was designed for the problems of the industrial economy, is failing to counter these harms. Also, and more generally, the administrative state is ill-equipped to define and assert public values and priorities relating to the design, implementation, and operation of digital architectures, systems, and processes. This preliminary concept paper is part of a project to reinvent the administrative state so that it is capable of governing the information economy effectively in ways that prioritize public accountability and oversight.

This concept paper imagines how to equip the administrative state to monitor the information economy and enforce mandates relating to monitoring and information production. To engage in real, meaningful oversight of information-economy, data-driven and algorithmically-driven activities, agencies need new monitoring tools and related enforcement capabilities. Some of the suggestions in this document have the goal of facilitating meaningful compliance with the kinds

[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, Chaz Arnett, Solon Barocas, Hannah Bloch-Wehba, Orly Lobel, Laura Moy, Gabrielle Rejouis, Christian Sandvig, Alicia Solow-Niederman, Michael Veale, Ari Waldman, and Lauren Willis for their participation in the academic workshop and their constructive and generous feedback.

of public mandates that exist, or might soon exist, now. Others are designed to equip regulators to understand the operation of digital architectures, systems, and processes more generally, so that they and/or Congress can determine how to structure new, more effective public mandates.

We first describe the kinds of information that regulators need and then consider the mechanisms that must be strengthened or created to ensure access to all of the necessary information. Next, we sketch a set of institutional changes designed to enable regulators, auditors, and the public to acquire, verify, and understand information about digital architectures, services, and processes. Finally, we consider steps that policymakers can take to help develop and deepen the pool of people with the technical and organizational skills required to conduct, audit, and oversee monitoring of digital architectures, systems, and processes. This document is very lightly footnoted but is accompanied by a curated bibliography of useful secondary resources.

1. Information Needed from Regulated Entities

We begin with an overview of the kinds of information that regulators need from regulated entities to discharge their mandates effectively.

Technical Information. It is important that agencies and auditors be explicitly empowered to request and, if necessary, compel access to all of the different kinds of information they need to oversee information-economy processes effectively. Some kinds of information that agencies need are relatively easy to describe and obtain using their traditional authorities. For example, when investigating discriminatory practices in advertising for housing or employment, the relevant agency needs to know at least what categories advertisers can use to target consumers. Digital architectures, systems, and processes, however, have grown increasingly complex, disaggregated, and emergent. The rise and dominance of complex machine learning systems raises specific issues, discussed below. Increasingly, therefore, all agencies need additional information in order to fulfill their public mandates.

Digital processes involve many levels of complexity. They involve various kinds of inputs (including code, data, engineering parameters) and produce various kinds of outputs (including both discrete results in particular cases and larger effects on populations or systems). They may operate over great distances and at very large scales (as with distributed networks of sensors or software developer kits designed for data collection). They may incorporate machine learning decisionmaking that resists explanation (as with both real-time ad placement based on population data and large language models based on web and social media inputs). Additionally, both inputs to and outputs of digital processes may change continually in ways that make external evaluation based on static reports impossible (as with the above examples and also with ride sharing pricing)

There are at least three categories of methods that agencies may need to employ to monitor compliance with public mandates, depending on the nature of the system under scrutiny and the nature of the law or regulation being applied.

Category One: Probing Black Boxes. Sometimes, agencies can execute their information gathering function by doing no more than probing private systems that happen to be connected to publicly accessible networks, sending test inputs and observing the resulting outputs (and other changes in behavior or state) to detect problems such as unlawful bias, regulatory noncompliance, or behavior inconsistent with public commitments. Attendees at our convening pointed to testing for discriminatory advertising on social media platforms as an information gathering imperative that can largely be accomplished in this manner, provided that the platforms are obligated to facilitate such testing (a matter to which we return in Parts 2 and 4, below).

Category Two: Ensuring Thorough and Responsive Disclosures. In many cases, however, outside black box testing alone cannot reveal all of the information needed to assess an information system's risks and harms. For example, if the question is not merely that a system may be producing discriminatory results but instead how the discrimination is produced and how to intervene, it will be essential to open the black box, specifying what information companies need to generate, store, and disclose to agencies (and auditors, as discussed more fully in Parts 2 and 4, below). As another example, regulators evaluating algorithmic strategies for risk prediction and mitigation might need to open the black box to assess whether the designers of such strategies have made appropriate allowances for, e.g., "black swan" events (as in the case of financial risk modeling), shifting patterns of baseline risk (as in the case of climate risk modeling), and cascading risks (as in the case of financial stability risks resulting from climate events.[2]

The precise information a company must generate, store, and disclose will vary by context, taking account of the goals of the operative laws and regulations, the kind of technology under scrutiny, the relevant history of the actors involved, and more. It is important not to be bogged down by increasingly arbitrary and rapidly changing distinctions between code, data, environment, etc. The important point is that agencies (and auditors, as discussed more fully in Parts 2 and 4, below) need access to the full range of information necessary to assess the design and performance of digital systems and processes and the actions and motivations underlying specific corporate behavior.

Access to source code will be valuable in some, but not all contexts. Sometimes disclosing code will be necessary; in other cases, code might be a distraction from more important questions having to do with data, training parameters, and other factors that structure the behaviors and outputs of digital systems and processes.

Where human programmers design systems that make decisions based on the machine representation of human-driven logic (as opposed to machine learning), agencies might insist on access to all of the underlying code. This will be especially true of legacy systems built before the recent expansion of machine learning decision making tools. Attendees at our convening pointed to probabilistic genotyping–techniques used to link genetic information left at a crime scene to stored genome information–as a field that still relies on human-coded logic rather than machine learning techniques. Reports suggest fields such as medical diagnosis, cybersecurity, and financial services are still relying on legacy expert systems, although we imagine many expert systems will be replaced by automated systems soon.[3]

With machine learning systems, in contrast, an agency might instead require access to the underlying weights, data, tools, techniques, and parameters used to develop (train and test) a model, as well as the tools and techniques used to assess the model and records of models trained but not deployed (for instance, records of all the A/B testing conducted on the model and internal records of decisions about why one model was prioritized or discarded). Sometimes, agencies may want access to information in its original format and/or to on-site systems and servers. At other

[2] Hilary Allen, *Regulatory Managerialism and Inaction: A Case Study of Bank Regulation and Climate Change*, J. L. CONTEMP. PROBS (forthcoming).

[3] Simon Preis, *Are Expert Systems Dead?*, TOWARDS DATA SCIENCE (Mar. 16, 2023), https://towardsdatascience.com/are-expert-systems-dead-87c8d6c26474.

times, they might require information to be produced in another format that they deem appropriate. Additionally, agencies may need to understand the architectures of systems used for data collection, exchange, querying, and content provision. This includes both information that is collected directly from or provided to entities in first party relationships and information collected in other ways, including via software developer kits incorporating APIs.

Because information-economy actors sometimes represent that they do not have particular kinds of information or elect not to undertake particular studies of their own systems, agencies also should be explicitly empowered to require companies to create and produce additional information.[4] For example, to support an investigation into whether a company's user interface amounts to a dark pattern, an agency should be able to instruct a company to run a series of A/B tests to try to rigorously demonstrate the effect of particular design choices on the company's own platform.

In some cases, regulators may want to observe systems operating in real time. For example, during election periods, it might be necessary to monitor the spread of illegal misinformation about polling places and times and about options for voting by mail.

Category Three: Full Reproducibility. In situations requiring the highest levels of scrutiny and oversight, agencies might demand *full reproducibility*, meaning replica snapshots of all of the data, code, and operating environment necessary to replicate the machine learning training, testing, and deployment steps that have been taken or to replicate the exact same outputs of the running model for given inputs. This might be necessary when the stakes for human wellbeing are especially high, meaning the cost of mistakes might be dire for the health or safety of individuals or the public or for critical public systems. Examples include voting machines, systems for managing power grids, and systems used to predict the quality of the food or water supply. Full reproducibility might also be justified after companies have been shown to have committed prior acts of fraud or gaming against a government monitor, a probationary remedy to detect and prevent repeat offenses. For example, if a company is shown to evade mandatory audits of algorithmic systems, it may be required to satisfy full reproducibility for new systems.

Meeting full reproducibility requirements may necessitate development of new tools and techniques that don't exist today, along with accompanying standards of documentation and organization. It may also preclude particular engineering techniques or design approaches that make full reproducibility impossible or difficult.

Importantly, in the case of very large and complex systems, preserving the possibility of full reproducibility also may require considerable resources (including storage, labor, time, and processing power), with associated costs both for companies and for the environment more generally. Thus, regulators wanting to require full reproducibility will need to think carefully about whether and when this approach would be warranted.

Non-Technical Information. The organizational structures of information economy companies are often fluid, and this has the side effect of frustrating efforts to understand, monitor, and ultimately govern the inner workings of these companies. If companies state that they are unable to provide the kinds of information described below because roles are not well-defined or

[4] Rebecca S. Eisenberg, *The Role of the FDA in Innovation Policy*, 13 MICH. TELECOMM. & TECH. L. REV. 345, 370 (2007) (arguing that the FDA's role includes both disseminating important information to the public and encouraging the production of new information that corporate actors might otherwise not be interested in producing).

relationships with stakeholders are informal, regulators might need to mandate more formal record-keeping arrangements. We return to this possibility in Part 3 below.

Organizational Structure and Lines of Accountability. Agencies need information about companies' internal organization to exercise effective oversight, an understanding of workflows that provide comprehensive assessment of accountability: an accountability graph. Organizational charts can sometimes be helpful, particularly for organizational design intended to grow an organization toward a particular goal. Organization charts provide a "snapshot" of an organization that communicates relationships, divisions of labor, and management arrangements. However, organizational charts have fallen out of vogue and any existing organizational chart might be insufficient when roles change quickly or do not adequately reflect the reality of the company. If a company is not utilizing or maintaining organizational charts, regulators need maps of workflows for relevant efforts in the company, including how work moves from one individual or team to another, how it is assessed, and when a decision by one team overrides a decision by another team. Working with auditors, organizations can transform these workflows into accountability graphs that must be properly maintained and updated as workflows change.

Internal Communications. As relevant to mapping and understanding the lines of accountability (both for particular decisions and for more general decisions about company policy), agencies will also need information about communications within the company. For example, regulators might want to know how teams in charge of governmental relationships or digital advertising sales influence the work of other parts of the firm. Understanding how different teams weigh in on decisions will help regulators understand how companies balance competing considerations and how they make final decisions.

Internal Policies, Training Materials, and Assessment Metrics. Agencies need information about the management of employees and divisions. Internal policies of concern to regulators might include product guidelines and methods for raising concerns about products, the contents and timing of employee orientations and ongoing trainings, and the metrics the company uses to evaluate job performance. Such indicators can help agencies understand the structure of incentives and influences that drive different sectors of the company.

External Communications. Communications with consumers and/or competitors are often relevant to assessing compliance with public mandates. For some information-economy actors, the volumes of such communications may be very large (for example, social media companies might communicate with millions of users on a daily basis), and different kinds of communications also may be directed to different audiences. Agencies will need information about these communications in formats that they can parse and analyze.

Companies might also need to disclose the nature and extent of their contacts with third parties that have or might have power to influence company policies, including government officials and other important groups of external stakeholders, such as angel and venture capital investors. The goal is for the agencies to understand how different sectors of the company interact with these actors (for example, what are the interrelationships among the staff managing government relations, those handling compliance, and those setting privacy policy?).[5]

[5] Evelyn Douek, *Content Moderation as Systems Thinking*, 136 HARV. L. REV. 526 (2022) (discussing the forms of internal organization that are necessary to enable public oversight and accountability).

2. Mechanisms for Obtaining Information from and About Regulated Entities

This section outlines the mechanisms through which regulators and other relevant actors might obtain the information they need to fulfill their duties effectively. Although the precise extent of authority varies from agency to agency, contemporary regulators already wield considerable monitoring authority.[6] In important respects, however, that authority still is not optimized to information-economy architectures, systems, and processes. Regulators also need better mechanisms for listening to the communities affected by information-economy actors.

The Basic Regulatory Monitoring Toolkit. Basic elements of the regulatory monitoring toolkit include the following. Generally speaking, existing agencies have and use these tools to varying extents.

Periodic On-Site Inspections. An agency might inspect a regulated entity's production process or outputs. Such inspections are common in industries that produce food or medicines for human consumption, although the efficacy of current inspection regimes is disputed.[7] As another example, the CFPB can use its supervisory authority to conduct on-site inspection of companies in the consumer finance industry. In the case of large depository institutions and affiliates, examinations are coordinated with prudential regulators and state regulators to ensure consistency with statutory requirements. During examination, CFPB examiners go on-site to observe, conduct interviews, review additional documents and information, transaction test, and assess compliance management.[8]

Periodic Submission of Information for Public Release and/or Pre-Approval. An agency might require regulated entities to disclose certain information to create a public record of their activities or, in some cases, for approval of those activities. For example, the SEC requires publicly traded companies to submit and make available to current and potential investors regular financial reports covering various matters. The CFPB requires covered financial institutions to report data on different types of services, including credit provision to small businesses, mortgage lending, and financial products such as credit cards and prepaid accounts. The EPA requires car companies to submit information indicating that they are complying with emissions requirements.

Periodic Audits and/or Supervision to Verify Compliance with Public Mandates. An agency might require regulated entities to submit to periodic audits or to regular supervision. For example, pursuant to the Health Information Technology for Economic and Clinical Health (HITECH) Act, the Department of Health and Human Services periodically audits covered entities in the healthcare industry to assess their compliance with the HIPAA Security Rule.[9] In the consumer finance and banking industries, some especially dominant firms have permanent auditor teams or supervisory

[6] Rory Van Loo, *The Missing Regulatory State: Monitoring Businesses in an Age of Surveillance*, 72 VAND. L. REV. 1563 (2019).

[7] Konstantinos Kotsanopoulos & Ioannis Arvanitoyannis, *The Role of Auditing, Food Safety, and Food Quality Standards in the Food Industry: A Review*, COMPREHENSIVE REV. FOOD SCI. & FOOD SAFETY (Aug. 3, 2017). The Centers for Medicare and Medicaid Services carries out on-site inspections of centralized medical testing facilities, providing another example of this model. *CLIA Program & Medicare Lab Services*, CTR. MEDICARE & MEDICAID SERVICES (Dec. 2021).

[8] Lorelei Salas, *What new supervised institutions need to know about working with the CFPB*, CONSUMER FIN. PROT. BUREAU (Jan. 9, 2023), https://www.consumerfinance.gov/about-us/blog/what-new-supervised-institutions-need-to-know-about-working-with-the-cfpb/

[9] Department of Health and Human Services Office for Civil Rights, *2016-2017 HIPAA Audits Industry Report* (Dec. 2020), https://www.hhs.gov/hipaa/for-professionals/compliance-enforcement/audit/index.html.

teams on site. The Federal Reserve carries out on-site inspections of supervised bank holding companies. The frequency and scope of the audit varies with the number of total consolidated assets and the size and complexity of the business.[10] The CFPB also exercises supervisory authority over both depository and non-depository financial institutions, prioritizing institutions subject to an assessment of risk. It requires reporting of certain kinds of information, requires supervised entities to generate and provide or retain certain records, and also spends time on-site at the offices and operation centers of supervised entities.[11]

Obligations to Create, Store, and Organize Information. An agency might require regulated entities to produce some forms of information, store it adequately, and provide access to regulators and/or auditors when requested, even though it might not be submitted periodically for review. For example, privacy consent decrees crafted by the FTC typically require the companies they cover to create and maintain certain records regarding their handling of covered data.

Limitations of the Basic Toolkit. Broadly speaking, the basic regulatory monitoring toolkit suffers from four problems.

One problem is inconsistency and incomplete coverage. Different agencies tend to rely on different kinds of information and different kinds of monitoring authority–a state of affairs that made sense when the sectors they regulated were more distinct but that makes far less sense now that digital systems, disaggregated data architectures, and data-driven, algorithmic processes have become common denominators. For example, the nature of the information needed to evaluate compliance with environmental requirements is different now that automotive systems are also digital systems. The nature of the information needed to evaluate compliance with anti-discrimination and consumer protection mandates has changed now that online advertising is served to end users via pattern-based, machine learning processes that ingest geolocation data collected by mobile communication, search, and social network providers. And no agency currently has clear authority to monitor the digital architectures, systems, and processes constructed and operated by platform entities whose operations span multiple economic sectors.

A second problem is that static, localized snapshots of information-economy processes that are networked and continually evolving are inadequate. As one example, a consumer protection or antitrust authority scrutinizing outputs of or inputs to the online advertising ecosystem must do more than simply order snapshots of activity from any one actor, even when that actor is a large advertising network or social media platform. The online advertising ecosystem is massive, distributed, and ever-changing. A snapshot-based approach will miss critical details occurring elsewhere in the ecosystem or before and after the snapshot. Part 1 offered several ways of thinking about the kinds of information that agencies will increasingly need to access on an ongoing basis.

Third and relatedly, the complexity of digital systems introduces new opportunities for gaming or evading regulatory mandates. As described in Part 1, black box techniques that interrogate the visible part of an agency's system from the outside can be detected from the inside, allowing companies to manufacture false, seemingly compliant data. Uber used analytics to detect and reject rides requested by city transportation authority employees and creators of computer viruses code

[10] *SR 13-21: Inspection Frequency and Scope Expectations for Bank Holding Companies and Savings and Loan Holding Companies that are Community Banking Organizations*, BD. GOVERNORS FED. RESERVE SYS. (Dec. 16, 2022), https://www.federalreserve.gov/supervisionreg/srletters/sr1321.htm.
[11] 12 U.S.C. §5514.

those viruses to behave differently when running inside a testing environment. Agencies must be able to verify that compliance is real.

A final problem involves accountability of the auditors and inspectors who function as vital monitoring intermediaries. Across the information economy, audit requirements and practices have become both increasingly widespread and increasingly controversial. Companies routinely develop internal policies to ensure compliance with public mandates, but in operation, those policies can become a series of checkboxes that do not fully reflect the relevant policy goals. In particular, when companies are in charge of translating public mandates into more granular obligations, they may do so in ways that reflect corporate preferences instead of public ones. In turn, auditors may equate internal policies with compliance rather than conducting more rigorous inspections of whether the organization's practices actually align with public mandates. When auditors work too closely with a company, there is also a risk that they will self-identify as part of the company's compliance team. We take these criticisms of audit and auditors very seriously. We also think, however, that attempting to eliminate audit and inspection requirements from the information-economy regulatory monitoring toolkit would be disastrous. And, as a practical matter, it is simply infeasible for even a substantially reinvigorated and adequate resourced administrative state to conduct all audits and inspections of all information-economy actors whose operations must be audited or inspected. The main question, then, is what we can learn from sectors where such requirements are extensive and where regulators and researchers have already identified and attempted to understand and address deficiencies. We return to this question in Part 4 below.

An Expanded Regulatory Monitoring Toolkit for the Information Economy. To help address the problems of incomplete coverage, emergent ordering and regulatory evasion, agencies need the following capabilities:

Universal Basic Toolkit. As a baseline, all agencies tasked with overseeing information-economy activity need all of the basic authorities described above in the first instance–i.e., without needing to wait until an investigation is opened or a consent decree is entered. Additionally, periodic reporting, on-site inspection, and audit or supervision requirements should be extended to broader sets of important information-economy actors. We return to questions surrounding the implementation of these requirements in Part 4, below.

Capacity to Ask Questions and Get Answers. Over and above requirements for periodic inspections and production of pre-determined forms of information, regulators and auditors also need to be in dialogue with corporate actors. For example, if data produced by the companies is in a format that does not allow adequate evaluation, regulators and auditors need to be able to require it in a different format. Regulators and auditors also need to be able to request additional information to understand the various matters described in Part 1, above, including, but not limited to: how architectures, systems, and processes for data collection and exchange are structured and operated, how training data or optimization parameters are chosen, what measures have been taken to de-bias training data, what kind of testing the company performs on its own systems and what results these tests have yielded, and the internal organization and accountability structures of regulated entities.

Capacity to Run Experiments, Test Capabilities, and Probe Black Boxes. As described more fully in Part 1, above, regimes mandating periodic disclosures generally will be insufficient to enable regulators to evaluate the outcomes of digital systems and processes. Just as companies regularly run experiments to determine engagement with different interface arrangements and different types of content, so regulators and auditors need to be able to run experiments or tests to determine whether other values are being protected. Such experiments and tests may be more effective when done with full system access. For instance, an agency might test a chatbot by asking it questions, as users might do, to determine whether it will provide manipulative voting information, but it can run those tests more effectively, and experiment with parameters and possibilities, from within the company's system. It is also essential that companies be forbidden from evading, discouraging, or corrupting black box testing. For example, companies may need to be able to provide unencrypted and non-proprietary access to systems. More proactively, companies should be required to configure their systems to increase the efficacy of black box testing. For example, a social media platform may be required to create and document a private API to provide regulators and auditors with programmatic access to public messages on its platform.

Authority to Tailor Monitoring Programmatically for Particular Sectors, Activities, and/or Systemically Important Actors. Agencies need the flexibility to determine which types of oversight are appropriate for which processes. To ensure that agencies themselves are accountable for the ways they use (or refrain from using) their monitoring authority, agency decisions about programmatic monitoring should be publicly disclosed and explained, and should be accompanied by official requests for information to help agencies assess the efficacy of their choices.

Additionally, there is need for an entity with authority to monitor certain operations of systemically important platform entities. We return to this issue in Part 4, below.

The Role of Publics in Regulatory Monitoring: Communities, Civil Society, Journalists, Academic Researchers, and Workers. Affected publics and more specialized organizations and groups have specific forms of knowledge and expertise that can assist in monitoring the information economy.

Involving Communities and Community Organizations in Oversight Activities. Public participation is essential for effective governance, as is harnessing the various forms of more specialized community expertise. In particular, communities are expert in the specific ways that digital technologies and processes harm or otherwise affect them. In a future concept paper, we will focus on mechanisms for including publics in all aspects of policy making and implementation. Here, we focus specifically on three useful mechanisms for inclusion in regulatory monitoring: community-led participatory research, participatory audits, and public advocates.[12]

In community-led research, members of the community identify priorities and needs, and professional researchers, if they are involved, take direction from community members and organizations. Research can inform policy development and program creation, and it can also help identify enforcement failures or priorities. For example, understanding how community members find or struggle to find rental housing might support a decision to subject online rental services and/or related advertising to higher scrutiny. Community-led research often involves working closely with local organizations or helping such organizations to form where they do not already exist. In turn, local organizations might need resources, logistical support, and adequate training to enable their meaningful participation. The partnerships between communities and professional

[12] Ben Palmquist, *Equity, Participation, and Power: Achieving Health Justice Through Deep Democracy*, 48 J. L. MED. & ETHICS 393 (2020).

researchers should aim at developing trust, sustained dialogue, and civic infrastructure for stakeholder participation.

Participatory audits involve publics in multiple ways. They incorporate interviews with users, consumers, workers, and members of affected communities. For example, auditing a ridesharing company should include engaging in dialogue with drivers and users to identify areas of concern. Affected publics and local organizations serving them can also help establish metrics and criteria for auditing.[13]

Public advocates are independent monitoring offices that receive complaints from members of the public and advocate on the public's behalf. Although their main function is to provide direct services, they should also be empowered to initiate investigations and produce reports to feed back into policymaking. Their direct contact with communities is an opportunity to inform agencies about the need to monitor certain sectors more closely or suggest enforcement proceedings.

Civil Society Organizations, Journalists, and Academic Researchers. Civil society organizations, journalists, and academic researchers have many kinds of specialized expertise relevant to seeing and understanding the information economy, and they produce essential research that can help to inform regulators' assessments of public harms and interests. The section in Part 4, below, on the proposed Public Research Institute discusses how these actors might get access to the data they need.

Whistleblowers. Whistleblowers play an important role in monitoring wrongdoing, but in order to perform that role more consistently and effectively, they require legal shelter. Today, whistleblowers most often rely on the protections offered by the National Labor Relations Act and the Dodd-Frank Act, which cover relatively narrow sets of wrongdoing and are not well matched to the kinds of wrongdoing that tech industry whistleblowers have revealed. In particular, whistleblowers should have legal protections when they report that companies have grossly misrepresented their activities to the public or to regulators. Employees should also be protected when they report that a company does not have adequate systems in place to deliver on its voluntarily assumed responsibilities, even when the employee does not have enough information to show that the firm is in fact not delivering on its promises. There should be explicit protection for whistleblowers who reveal deliberate withholding or misrepresentation of information required by agencies and necessary for effective oversight.

3. Implementation Mechanisms

Some of the capabilities outlined in this concept paper can be implemented or partly implemented using existing authorities, but others may require expanded authorities and/or new mechanisms. In general, we think that important benefits will flow from more explicitly defining new monitoring capabilities of administrative agencies in relation to information-economy activities. Resources now spent litigating the scope of administrative authority would be better used developing capabilities appropriate to the information economy and learning how to use them effectively.

[13] OUR DATA BODIES PROJECT, https://www.odbproject.org (promoting the work with local organizations to design practices for collecting, storing, and sharing data about communities).

Broad, Forward-Looking Regulatory Monitoring Authority. Some agencies currently have authority to conduct broad, forward-looking regulatory monitoring; we think all agencies that regulate information-economy actors should be similarly empowered to do so. So, for example, under §6(b) of the FTC Act, the FTC can conduct studies even if they do not have a defined law enforcement objective. Additionally, it can require an entity to file "annual or special . . . reports or answers in writing to specific questions" to provide information about the entity's "organization, business, conduct, practices, management, and relation to other corporations, partnerships, and individuals."[14] Similarly, under sections 1024-26 of the Consumer Financial Protection Act of 2010 (Title X of the Dodd-Frank Act), the CFPB has broad forward-looking monitoring and supervisory authority, which it may exercise for the purpose of: "(A) assessing compliance with Federal consumer financial law; (B) obtaining information about a supervised institution's activities and compliance systems and procedures; and (C) detecting and assessing risks to consumers and to markets for consumer financial products and services."[15]

Critically, agencies should be empowered to conduct such inquiries using all of the mechanisms for obtaining information in the expanded regulatory monitoring toolkit described in Part 2, above. (We note, as well, that there are other important differences in the ways these authorities are currently designed. As one example, although the FTC has broad investigative authority, it currently must satisfy strict evidentiary requirements before engaging in rulemaking. We will consider rulemakings and other policy mechanisms in a future concept paper.) The information gained through these monitoring activities should be available without limitation to enforcement staff.[16] Additionally, there is a need to rebalance current limits on public disclosure of such information to facilitate greater accountability to customers and the general public; we discuss this issue in Part 4, below.[17]

Express Authority to Promulgate Standards for Regulatory Monitoring and Auditing or Supervision of Digitally-Mediated Activities. As they use the expanded regulatory monitoring toolkit described in Part 2, above, regulators may identify standards that facilitate monitoring, supervising, and auditing activities across the information economy. For example, they might find that certain forms of internal organization make it easier to understand how to direct questions, or that a specific format for storing and producing data facilitates longitudinal (intra-company) and/or cross-company comparison of the results of audits. In such cases, regulators need to be able to set industry-wide standards without the need of negotiating or imposing specific duties one company at a time. Ideally, this process would be conducted or coordinated by the new hub entity that we describe in Part 4, below.

Express Authority to Prescribe Best Practices in the Design of Digital Architectures, Systems, and Processes to Enable Regulatory Monitoring. Digital architectures, systems and processes need to be designed in such a way as to make compliance with public mandates verifiable. In the case of the FTC, promulgation of best practices is now achieved chiefly via consent decrees. So, for example, in the context of an order dealing with unfair and deceptive practices related to collection

[14] 15 U.S.C. §§ 46.

[15] 12 U.S.C. §5514.

[16] *30 Supervisory Highlights*, CONSUMER FIN. PROT. BUREAU 1, 35 (Summer 2023) (showing how the CFPB's supervisory activities can result in and support public enforcement actions).

[17] Peter Conti Brown, *The Curse of Confidential Supervisory Information*, BROOKINGS (Dec. 20, 2019), https://www.brookings.edu/articles/the-curse-of-confidential-supervisory-information/ (describing how regulators and Congress might relax the rules shielding bank supervisory information from public disclosure to improve accountability of the financial system without impairing the deliberations between banks and bank supervisors).

of location data, the FTC might specify details regarding the design of the consent interface and/or prohibit certain design choices, and it might mandate regular data deletion schedules. Such efforts are a good start, but we think the FTC (or any other agency) should not need to wait until the consent decree stage before defining and prescribing best practice obligations. We also think that conceptions of the kinds of best practices that are relevant can and should be recalibrated to enable meaningful oversight. So, for example, the FTC (or the new hub entity described in Part 4, below) could require a platform company to maintain logs of all app developers that have installed software developer kits (SDKs) that collect and transmit precise geolocation information.

Hard Limits on Law Enforcement and/or National Security Access to Information Collected through Regulatory Monitoring. Broadened regulatory monitoring authority requires correspondingly more effective safeguards to prevent the information from spilling over into unrelated criminal, immigration, and national security investigations. Currently, statutes regulating information collection tend to include fairly flexible exceptions benefiting such investigations.[18] One notable exception is the Census Act, which strictly and specifically prohibits the "use . . . for any other purpose other than the statistical purpose for which [information] is supplied." 13 U.S.C. § 9(a)(1). As the scope and complexity of information collection by both private and government entities continues to grow, we think that comparably strict standards should shield information collected through regulatory monitoring against access for unrelated purposes, and that law enforcement and national security investigators should follow separate, well-defined processes to gain access to information to meet their legitimate needs.

Authority to Enforce Compliance with Regulatory Monitoring-Related Obligations. Compliance with regulatory monitoring is essential for meaningful enforcement of public mandates. Currently, many information-economy actors flout the information production obligations that regulators attempt to impose. This dynamic is especially pronounced where the largest and most powerful firms are concerned, and some of those firms also have stonewalled in the face of information requests from Congress and/or courts.

When regulated entities fail to comply with regulatory monitoring obligations, agencies should be empowered to impose (and enforce payment of) significant fines. Fines should scale in a way that is commensurate with company size and should ascend for repeat violations. They should encompass the full spectrum of regulatory monitoring obligations, including disclosure obligations, obligations to provide access to regulators and auditors, and obligations to respond to inquiries and facilitate experiments. To help ensure that fines represent meaningful deterrents, some types of fines should accrue automatically. Agencies should publish schedules of these fines along with information about how they adjust for company size. To help ensure that both agencies' authority and its own policymaking authority are respected, Congress should augment agencies' budgets by adequate amounts specifically earmarked for enforcement against regulatory monitoring violations. For more severe violations, a more extensive menu of sanctions might include company reorganizations and reorganizations of disaggregated data architectures, as needed to enable more effective monitoring of compliance with public mandates.

Agencies also should have express authority to hold certain important actors individually accountable for their companies' failures to comply with regulatory monitoring obligations. At minimum, civil fines should apply to executives who own shares giving them 50% or more of

[18] Erin Murphy, *The Politics of Privacy in the Criminal Justice System: Information Disclosure, the Fourth Amendment, and Statutory Law Enforcement Exemptions*, 111 MICH. L. REV. 485 (2013).

shareholder voting power. Additionally, we think that it is worth considering whether certain repeated regulatory monitoring violations should trigger civil fines for all top executives and for board members (in the case of public companies) or major investors (in the case of nonpublic companies).[19] In both cases, the authorizing language should specify that fines assessed against individual executives shall not be insured against or indemnified.

Last but not least, regulated entities' internal organization must allow oversight. Both the FTC and the CFPB have sometimes used consent decrees to institute new structures for independent assessment and board-level reporting on compliance matters. We think that this approach holds promise and should be formalized and extended. When the internal organization of the firm is too complex to allow oversight, regulators should have express authority to order internal restructuring to create clear lines of accountability, as described in Part 1, above.

4. Institutional Design for Monitoring and Enforcement

Equipping administrative agencies with new capacities requires changes in the institutional design of the administrative state. This section proposes institutional reforms oriented toward empowering regulators to see and understand information-economy architectures, systems, and processes, empowering researchers who seek to study private sector digital systems and processes, and restructuring the relationships among auditors, firms, and regulators.

A particularly challenging question is whether it would make more sense to create a new agency dedicated to monitoring use of data-driven, algorithmic tools and processes across all realms of economic activity or whether it is preferable and/or necessary to equip all agencies with new resources. We think that this is not an either/or question and that a properly designed new entity can function as a central hub within a network of new digital monitoring and enforcement capabilities designed to mirror the structure of the information economy. Another question is whether it is necessary to redraw the jurisdictional mandates of existing agencies to account for the cross-cutting nature of certain information-economy activities. We take no position on the second question in this particular concept paper. Our proposal for a Digital Architectures, Systems, and Processes Oversight Board designed as a hub to support monitoring of digital processes is independent of the jurisdictional arrangements designed for other agencies. Additionally, we propose creation of two more narrowly scoped entities–a Public Research Institute and a Digital Processes Audit Oversight Board–both of which could be sited within the new hub.

Digital Architectures, Systems, and Processes Oversight Board. The federal government should create a new Digital Architectures, Systems, and Processes Oversight Board to support regulatory monitoring of digital architectures, systems, and processes. The new board would have a hybrid function. It would perform certain functions that are more effectively centralized and that are necessary for the administrative state, taken as a whole, to understand information-economy architectures, systems, and processes. Simultaneously, a hub-and-spoke model for collaboration between the new board and existing, domain-specific agencies would facilitate use and iterative improvement of knowledge and techniques developed in the hub. This overall structure would

[19] *FTC Takes Action Against Drizly and its CEO James Cory Rellas for Security Failures that Exposed Data of 2.5 Million Consumers*, FED. TRADE COMM'N. (Oct. 24, 2022), https://www.ftc.gov/news-events/news/press-releases/2022/10/ftc-takes-action-against-drizly-its-ceo-james-cory-rellas-security-failures-exposed-data-25-million (binding CEO James Cory Rellas directly to specific data security requirements for his role in presiding over unlawful business practices).

facilitate improved understanding of the harms that need to be addressed, the specific ways that digital technologies underlie and contribute to those harms, the kinds of questions that regulators need to ask firms, and the kinds of information they need to have produced.

The existing administrative state includes many examples of hub-and-spoke experiments, some more successful than others. Some involve centralized oversight and coordination of policy. For example: the Office of Information and Regulatory Affairs was created to enable centralized review and cost-and-benefit analysis of proposed regulatory initiatives; the Office of the Director of National Intelligence was created to coordinate knowledge-sharing within the intelligence community; and the Financial Stability Oversight Council was created to facilitate assessment of systemic financial risk and coordinate corrective measures.

Other examples of hub-and-spoke models have involved centralized provision of technical and research expertise. For example, the Chief Statistician in the Office of Management and Budget coordinates the activities of the various U.S. federal statistical agencies and helps them to work closely with existing, domain-specific agencies and departments. The Advanced Research Projects Agency (later renamed Defense Advanced Research Projects Agency) was established during the Cold War to advance cutting-edge scientific and technology research; its projects included the predecessors of the internet and the global positioning system. In 2007, Congress established a new ARPA-E within the Department of Energy to pursue advanced energy-related research. The 18F consultancy within the Government Services Administration and the more recently created US Digital Service helps agencies build digital infrastructures for the provision of public services.

The hub-and-spoke arrangement that we envision would combine elements from both kinds of models, though it would look more like the models in the second group. The new DASPOB would be tasked with performing three kinds of functions, with the possible addition of a fourth.

First, it would develop protocols, best practices, and technical expertise for regulatory monitoring of digital architectures, systems, and processes, and it would operate as a consultancy to supply other agencies with the resources they need to do their jobs – including additional, domain-specific monitoring and enforcement – effectively.

Second, , the DASPOB would be expressly empowered to elicit a wide range of general information about the operation of digital architectures, systems, and processes technologies and share it with existing agencies to use in fulfilling their more specific mandates. In particular, the DASPOB's regulatory monitoring authority would extend to the architectures, systems, and processes constructed and operated by platform entities whose operations span multiple economic sectors.

Third, for particular, systemically important platform entities, the DASPOB would be empowered to impose continuing, on-site supervision.[20] For example, in connection with its information collection function, the DASPOB might require search and social media companies to disclose certain information about their distributed architectures for data collection, or information about their optimization parameters.

[20] This terminology follows emerging consensus on the importance of training special kinds of scrutiny on actors with especially pervasive reach. In this document, we do not propose specific thresholds for determining when a firm qualifies as systemically important.

Fourth, and more speculatively, should Congress ultimately choose to enact new public mandates for systemically important platform entities, the DASPOB could oversee and enforce those mandates. That possibility, however, is beyond the scope of this preliminary concept paper.

Structuring a truly collaborative relationship between the DASPOB and existing agencies is key to both parts of this proposal. Each agency has domain knowledge that is fundamental to guide its own regulatory monitoring and enforcement activities. Each therefore might want to solicit different kinds of support from the hub. All would have interests in receiving at least some of the additional information the DASPOB would elicit and share. If done properly, centralizing certain functions relating to the development of capabilities for regulatory monitoring and to the provision of useful information can produce results that are additive rather than subtractive, encouraging information sharing and discouraging interagency turf battles and other forms of unproductive competition.

Public Research Institute. The federal government should create a Public Research Institute with a twofold mission.

First, the Public Research Institute would ensure access to private sector information for independent academic researchers, journalists, and civil society researchers. Researchers seeking to study the societal impacts of the information economy have experienced difficulties gaining access to information about digital architectures, systems, and processes, and some have been sued after gaining access in ways not sanctioned by the companies whose operations they sought to study. One challenge is that different research projects need different data; another is that some projects can be run remotely while others may require the ability to observe processes or run experiments in the original environment. In some cases, qualitative research might require participant observation or interviews with company's employees. The Public Research Institute would establish criteria for gaining access and administer access requests.

Second, the Public Research Institute would manage data issues raised by research projects and company disclosures. It would develop protocols for evaluating and managing secrecy, privacy, and/or security risks potentially raised by independent or public research projects and, relatedly, for evaluating and managing secrecy, privacy, and/or security objections raised by companies to disclosure of information about their operations.

Often, information economy actors will argue that they cannot or ought not disclose information due to concerns about trade secrecy, user privacy, or data and system security, or because revealing too much will permit gaming or hamper law enforcement, among other reasons. These concerns are important and worth acknowledging, but their assertion in particular contexts may seem overbroad or pretextual. Moreover, concerns about trade secrecy, privacy, and/or security should not be permitted to frustrate effective public oversight. The Public Research Institute can be tasked with developing procedures for managing the risks of disclosure while enabling disclosure to go forward. For example, it could develop protocols for sharing information with partners and researchers in secure disclosure environments that build on the Federal Statistical Research Data Center (FSRDC) model, within which researchers are subject to controls on the ways they are permitted to access and disseminate covered information, or on existing models used by researchers for sharing medical data.[21] Alternatively, it might permit companies to utilize

[21] Christopher J. Morten, Gabriel Nicholas & Salomé Viljoen, *Researcher Access to Social Media Data: Lessons from Clinical Trial Data Sharing*, 38 BERKELEY TECH. L.J. (forthcoming 2024) (discussing useful models of sharing medical data).

techniques such as differential privacy or synthetic data to release data while reducing attendant risks. As another option, in some circumstances, companies might be permitted to use techniques such as zero knowledge proofs, secure multiparty computation, and cryptographic commitments to prove certain system attributes without revealing additional information.

Digital Processes Audit Oversight Board. The activities and outputs of auditors, supervisors, and other third party compliance intermediaries must be subject to stricter oversight. Auditors are critical actors in the monitoring and enforcement landscape but, for the most part, have not been the sustained focus of thinking about regulatory reform. One notable and relatively recent exception is the chain of events leading up to the creation of the Private Company Accounting Oversight Board (PCAOB) for the financial sector. Another notable exception is the system for banking supervision, which is a term used to capture a broader process of continuous oversight that functions in addition to periodic audits and has developed over several decades through a number of iterations. In what follows, we borrow to an extent from those examples but also recognize that the particular skill sets required for oversight of digital processes demand a somewhat different approach.

The federal government should create a new Digital Processes Audit Oversight Board (DPAOB). The DPAOB would be responsible for the independent and public oversight of auditors and supervisors for digital systems and processes. It would set standards for conducting audits and ongoing supervisory processes and for certifying and reviewing the results of such processes. Such standard-setting is not without risks. The standards themselves might not be optimal, and auditors might become more interested in avoiding DPAOB's inspection than actually producing high-quality audits. To help offset those risks, the DPAOB would also set standards for training, certification, and discipline of auditors and supervisors. Evidence from studies of financial auditors indicates that inspection of auditors against industry standards tends to improve the quality of audits and make audits more easily readable and comparable.[22]

The DPAOB would gather information from existing agencies about their needs and experiences, and it would provide support to existing agencies wishing to supplement DPAOB standards with additional standards tailored to their particular missions and needs. It would receive information about and conduct preliminary investigations of violations of federal audit and supervisory standards. It would have authority to issue fines to and/or suspend the licenses of auditors who violate its standards. As appropriate, it would refer more severe violations to investigation and enforcement branches similar to any other regulatory violation.

The relationships among auditors/supervisors, the administrative state, and firms can be structured in different ways. In some cases, government employees should conduct the audits or lead audit/supervision teams, while in other cases external auditors may be more appropriate. When external auditors are used, a common problem has been that auditors are accountable to management, resulting in a conflict of interest and low-quality audits. In addition, experience with financial auditors teaches that auditors can be prone to various kinds of groupthink and as a result can miss–or deliberately overlook–warning signs at particular firms and within industries or systems more broadly. As a way of mitigating these problems, some have proposed that

[22] Daniel Goelzer, *Audit Oversight and Effectiveness: Understanding the Past and Looking Toward the Future*, CPA J. (2021), https://www.cpajournal.com/2021/05/25/icymi-audit-oversight-and-effectiveness/; Daniel Aobdia, *The Impact of the PCAOB Individual Engagement Inspection Process — Preliminary Evidence*, 93 ACCT. REV. 53 (2018); Takiah Iskandar, Ri a Sari, Zuraidah Mohd-Sausi & Rita Anugerah, *Enhancing auditors' performance: The importance of motivational factors and the mediation effect of effort*, 27 MANAGERIAL AUDITING J. 462 (2012).

shareholders participate in assigning auditors, but we think that auditors also should be accountable to the general public. To incentivize public accountability and enhance audit quality, regulatory agencies should assign auditors randomly and administer their compensation, which should be funded through fees paid by companies.[23] Additionally, the DPAOB should develop performance standards for auditors that encompass and reward the exercise of adversarial investigation and independent judgment.

5. Building the Pipeline

A final set of open questions concerns how the government can help enlarge the pool of individuals who possess the skills necessary to audit complex technical information about information economy processes. Like financial auditors, auditors for digital processes require specific and extensive skills; however, there is a robust pipeline for developing and honing financial accounting and audit skills and no comparable pipeline for acquiring the relevant skills for auditing digital processes.

To address this deficit, the federal government should invest in the development of curricula and training programs for auditors and supervisors for digital processes and should invest in the people entering such programs. Toward both of these ends, it could design and implement a large-scale public service program, similar to the public programs of the New Deal, to train and employ digital analysts. At minimum, it should offer grants to universities, community colleges, technical institutes, and other institutions interested in developing new programs or improving existing ones. Additionally, it should offer scholarships for prospective students and create fellowships and other research opportunities for more advanced study.

There are many examples of existing fellowship programs the federal government has created to build talent pipelines in areas where more and better trained professionals are needed. The Office of Energy Efficiency and Renewable Energy has implemented a series of fellowships designed to train scientific professionals and policy makers on subjects such as renewable energy and climate justice and on working with different stakeholders. The Cybersecurity Talent Initiative seeks to recruit and train a cybersecurity workforce by giving participants the opportunity to work for two years in an agency with relevant needs. The Dwight David Eisenhower Transportation Fellowship Program aims at building talent in the transportation sector.

In the specific case of digital audit, the programs created and supported by the federal government to build the pipeline for digital auditors should develop four sets of skills:

<u>Technical Skills</u>. Auditors need to understand how to interrogate and critically evaluate complex, data-driven digital systems and processes. University-level computer science and data science programs typically do not teach these skills, focusing instead on programming and optimization skills.

<u>Societal Impacts of Digital Systems</u>. Auditors for digital architectures, systems, and processes also need important kinds of non-technical knowledge. In particular, they need training to understand how users interact with digital technologies and how those technologies and the business models that shape their design, implementation, and use affect users, communities, and social institutions.

[23] David Khan, *Who's the Boss? Controlling Auditor Incentives Through Random Selection,* 53 EMORY L. J. 391 (2004); Patrick Hurley, Brian Mayhew & Kara Obermire, *Realigning Auditors' Accountability: Experimental Evidence,* 94 ACCT. REV. 233 (2019).

Qualitative Skills. To help ensure that compliance is not reduced to a meaningless checklist, auditors also should receive training in qualitative evaluation methods. At minimum, they should be able to inquire into the reasons for corporate behavior and to verify that organizations are sufficiently documenting those reasons.

Working with Stakeholders. As discussed above, auditors should collaborate with workers, users, consumers, and affected communities in order to gain an adequate understanding of the impacts of the technologies they are auditing. Like the programs developed by the Office of Energy Efficiency and Renewable Energy, the programs we envision should train auditors to work with different stakeholders.

6. Conclusion

Reinventing the tools that administrative agencies have at their disposal to monitor information-economy actors and activities is essential for recentering the public and public values in governance. Together, the proposals described here offer a blueprint to begin that process.

The proposals described here also represent only a first step toward the larger goal. As we noted at the outset, this preliminary concept paper is part of a larger project to reimagine the administrative state for the information era. Future modules will explore at least the following six additional issue clusters: (1) how government *builds and procures digital tools and systems*, (2) the *policy mechanisms* necessary to develop effective public mandates regarding information-economy actors and activities, (3) mechanisms for *meaningful inclusion of various publics* in information-economy governance, (4) mechanisms for *meaningful enforcement of public mandates*; (5) the *institutional design* of an administrative state optimized for the information era, and (6) *rule of law requirements* for governing information-economy architectures, systems, and processes. We expect the proposals in this document to evolve as work on the other modules proceeds.