Legal Priorities Project

Catastrophic Risk Review

Michael Livermore

LPP WORKING PAPER SERIES N° 3-2022

This paper was the second prize winner of the Legal Priorities Cost-Benefit Analysis Writing Competition

Catastrophic Risk Review

Michael Livermore

September 1, 2022

Summary

In the United States, cost-benefit analysis plays a substantial role in government policy making in traditional regulatory areas, such as automobile safety and air quality. But the management of catastrophic risk has largely fallen outside the purview of cost-benefit analysis. The primary reasons that cost-benefit analysis has played a less important role in the context of catastrophic risks are procedural, rather than methodological. Although improvements can be made to cost-benefit analysis techniques to better account for catastrophic risks, the more important set of reforms—and the ones discussed in this proposal—are institutional.

In the U.S. regulatory context, cost-benefit analysis is embedded in the process of regulatory review, which is almost entirely reactive. Agencies initiate regulatory proposals and then work with the Office of Information and Regulatory Affairs (OIRA) to analyze the effects of their proposals and make revisions in light of that analysis. OIRA has very little influence over agency agenda setting, which is instead dominated by legal and political factors that are largely unmoored from cost-benefit considerations.

For traditional regulatory areas, the reactive posture of OIRA may not be ideal, but it is not altogether debilitating. In these areas, robust administrative agencies with substantial statutory mandates and long traditions of regulation are relatively well-positioned to identify and respond to policy needs. OIRA's role of channeling and coordinating agency energies helps produce rules that are better justified and more likely to lead to net social benefits.

But the situation is very different for catastrophic risks. Regulatory review is well-suited for reining in overactive agencies or delaying or stopping imprudent agency action. But there are a wide range of catastrophic risks that are more likely to be exacerbated by government inaction than government action. Because OIRA does not evaluate agency agenda setting, cost-benefit analysis is not applied to the most consequential government decisions concerning catastrophic risks—the choice not to act.

The institutional reform discussed in this proposal is a **Catastrophic Risk Review** process, spearheaded by OIRA, that would examine catastrophic risks and potential government responses through a cost-benefit lens. This review process would build on earlier experiments in which OIRA has played a more proactive role in initiating regulatory actions. The two most successful of these experiments were the practice of prompt letters under the George W. Bush administration and the regulatory lookback undertaken by the Barack Obama administration. The Catastrophic Risk Review will also take advantage of OIRA's experience in cross-agency coordination and harmonization, including the Obama administration's interagency working group on the social cost of carbon, and the Bush administration's Circular A-4 guidance document on cost-benefit analysis methodology.

The purpose of this process would be to find areas where tangible, cost-benefit justified policy steps can be undertaken to manage catastrophic risks. The Catastrophic Risk Review would be overseen and partially staffed by OIRA, but should also include an interagency coordination group. There should be at least one round of public comments to help identify risks and potential responses. The Review should have a relatively sweeping purview, examining regulatory as well as other governance tools, such as research subsidies or even new legislation, that may be appropriate.

Ideally, Catastrophic Risk Review would lead to an ongoing process of identifying and addressing this important category of risk. Right now, there is, essentially, a backlog of unaddressed catastrophic risk, so a substantial initial effort is needed. After that backlog has been cleared, a more regularized updating process can evaluate existing efforts on identified risks and determine whether new risks merit further attention.

The products of a Catastrophic Risk Review could include new actions undertaken by executive agencies, recommendations to Congress for legislation or funding, or guidance on methodological updates to cost-benefit analysis to better account for catastrophic risks.

Proposed language

Executive Order 13,563 includes the following language on retrospective analysis, which can serve as a template for Catastrophic Risk Review:

> Sec. 6. Retrospective Analyses of Existing Rules. (a) To facilitate the periodic review of existing significant regulations, agencies shall consider how best to promote retrospective analysis of rules that may be outmoded, ineffective, insufficient, or excessively burdensome, and to modify, streamline, expand, or repeal them in accordance with what has been learned. Such retrospective analyses, including supporting data, should be released online whenever possible.
>
> (b) Within 120 days of the date of this order, each agency shall develop and submit to the Office of Information and Regulatory Affairs a preliminary plan, consistent with law and its resources and regulatory priorities, under which the agency will periodically review its existing significant regulations to determine whether any such regulations should be modified, streamlined, expanded, or repealed so as to make the agency's regulatory program more effective or less burdensome in achieving the regulatory objectives.

The proposed language, which could be included in a more general Executive Order or Presidential Memorandum on regulatory review, is as follows:

> Sec. X Catastrophic Risk Review (a) Catastrophic risks are those that threaten the continued viability of the national community. Risks that arise from pandemics, environmental crises, asteroids, weapons of mass destruction, or other advanced technologies (such as artificial intelligence and bioengineering) could profoundly disrupt the future of humankind. Identifying, analyzing, and addressing such risks are among the most critical tasks of the regulatory system.
>
> (b) The Office of Information and Regulatory Affairs shall convene a Catastrophic Risk Review Working Group to develop and initiate a government-wide assessment and review of catastrophic risks. Each agency, consistent with law and its resources and regulatory priorities, shall submit to this Working Group a preliminary plan to identify and examine catastrophic risks within its mandate. After opportunity for public comment and engagement, but no later than 300 days past the date of this order, the Working Group shall issue recommendations for regulatory or legislative reforms to facilitate consistent evaluation and management of catastrophic risks.

CONTENTS

INTRODUCTION

When Robert Oppenheimer was asked about his memories of the first successful test of a nuclear bomb at the Trinity site near Los Alamos, New Mexico, he recalled a line from the Bhagavad Gita where the god Vishnu takes on a terrifying, multiarmed form and says, "Now I am become death, the destroyer of worlds."[1] For many, the successful test of nuclear arms confirmed the human capacity to generate catastrophic risks—to become world-destroyers. As technological development has continued apace in the past three-quarters of a century, that potential has become only increasingly clear.

Fortunately, many catastrophic risks can, in theory, be efficiently managed. Although the risk of nuclear war is ever present, substantial steps have been taken to control access to bomb-making materials and to reduce existing stockpiles. At the same time, our understanding of the physical principles that underlie nuclear arms have also engendered the development of nuclear energy, a technology that, while it carries its own risks, also offers significant advantages as a carbon-free and stable source of electricity.

The challenge, as always, is to identify government interventions that generate risk-reduction benefits that justify their costs. In the United States, the use of cost-benefit analysis has a decades-long pedigree and is a well-entrenched part of the administrative state. However, evaluation of catastrophic risks using cost-benefit principles lags more traditional regulatory areas such as automobile safety and air quality. To bring catastrophic risks into the cost-benefit fold, reforms are needed.

Some of these reforms are methodological. When future costs and benefits are discounted, steps to avoid even extreme or catastrophic harms may be treated as having little value in present-day terms if the harms they avoid occur far enough in the future.[2] Catastrophes may also impose harms that are different in kind than standard risks. The harm of a catastrophe that wipes out all of humankind is not captured by the sum of the value of statistical life of each person alive at the time of the event because such a catastrophe would cut off the possibility of all future lives.[3] Catastrophic risks may involve some effects that are currently difficult to value and

---

1 KAI BIRD & MARTIN J. SHERWIN, AMERICAN PROMETHEUS: THE TRIUMPH AND TRAGEDY OF J. ROBERT OPPENHEIMER 309 (2006).

2 *See generally* CHRISTIAN GOLLIER, PRICING THE PLANTE'S FUTURE: THE ECONOMICS OF DISCOUNTING IN AN UNCERTAIN WORLD (2012); John Broome, *Discounting the Future* 23 PHILO. & PUB. AFFAIRS 128 (1994). For example, a trillion-dollar harm five hundred years from now, discounting to present value at a five percent discount rate, is a bit over twenty-five dollars.

3 *Cf.* TOBY ORD, THE PRECIPICE: EXISTENTIAL RISK AND THE FUTURE OF HUMANITY 52–53 (2020) (noting the "vast future that existential risks threaten to foreclose").

are left unquantified using existing techniques.[4] Methodological reforms of the standard practice of cost-benefit analysis in the United States could improve how such catastrophic risks are valued.

However, the primary barriers to appropriate treatment of catastrophic risks in U.S. regulatory decision making are not methodological. Even if appropriate methodological reforms were made, the institutional context of how cost-benefit analysis fits into the regulatory process would lead to inadequate attention to catastrophic risks. Although it is possible that a catastrophic risk could be created or exacerbated by an agency action of some kind, it is far more likely that the failure to act, rather than an action imprudently undertaken, would contribute to catastrophic risks. Indeed, agency inaction almost certainly already contributes to catastrophic risks on many fronts—climate change is an obvious example, but almost all categories of catastrophic risks, including those arising from advanced artificial intelligence, asteroid impacts, pandemics, weapons of mass destruction, and bioengineering, are exacerbated by the failure of U.S. policymakers to take steps that are currently available and that would be cost-benefit justified, even using current cost-benefit analysis methodologies.

For cost-benefit analysis to play a useful role in informing regulatory decision making regarding catastrophic risks, the primary reforms that are needed are institutional rather than methodological. Currently, cost-benefit analysis is used during a reactive process of regulatory oversight: agency agendas are developed on the basis of various legal, policy, and political criteria, and cost-benefit analyses are prepared of individual regulatory proposals. These analyses are then reviewed by the Office of Information and Regulatory Affairs (OIRA) in the White House and become part of the administrative record that is reviewed by courts in the course of litigation challenging agency decisions, typically under the Administrative Procedure Act (APA). Cost-benefit analysis is an important input during the process of regulatory design and analysis, but it plays a much more limited role in the context of agency agenda setting.[5] Given this institutional arrangement, it is extremely difficult for cost-benefit analysis to be used to call attention to underappreciated risks.

There have been attempts by prior administrations to leverage OIRA's cost-benefit expertise to inform efforts by the White House to direct agency agenda setting. The two most successful were the practice of prompt letters initiated by Administrator John Graham during the George W. Bush administration and the

---

[4] *See* Jonathan S. Masur & Eric Posner, *Unquantified Benefits and the Problem of Regulation under Uncertainty*, 102 CORNELL L. REV. 87 (2016) (examining degree to which agencies fail to quantify cost and benefits and possible explanations).

[5] *See* Michael A. Livermore & Richard L. Revesz, *Regulatory Review, Capture, and Agency Inaction*, 101 GEO. L. J. 1337 (2013).

regulatory lookback effort initiated by Administrator Cass Sunstein during the Barack Obama administration.[6] Prompt letters were used by OIRA to direct agency attention to cost-benefit-justified actions that the agencies were not undertaking. Some prompt letters were deregulatory in nature, and agencies were directed to existing regulatory interventions that OIRA believed should be either revised or rescinded.[7] But prompt letters were also used to address the failure of agencies to address certain risks. The regulatory lookback process was initiated to urge agencies to review their existing stock of regulatory requirements to identify rules that were good candidates for revision on cost-benefit criteria. In the lookback process, all agencies were directed to submit reports to OIRA on the results of this review process, and many agencies identified, and ultimately adopted, reforms that led to substantial cost savings.

Building on these prior efforts, this essay proposes a Catastrophic Risk Review process, spearheaded by OIRA, that would examine opportunities across the federal government to address catastrophic risks, and would also identify governance gaps where existing authority is inadequate. The purpose of this process would be to find areas where tangible, cost-benefit-justified policy steps could be undertaken to manage catastrophic risks. The potential policy measures examined should include regulatory interventions as well as other governance tools, including research subsidies, that may be appropriate. This process could be used to inform agency agenda setting as well as OIRA's annual Report to Congress on the Costs and Benefits of Federal Rulemaking.

Because OIRA is already severely over-taxed with its existing functions,[8] there is a concern that any additional responsibilities would necessarily require OIRA to shift personnel from its core functions. For this reason, were the administration to undertake a Catastrophic Risk Review, it should request additional resources from Congress to appropriately staff this initiative. OIRA is chronically underfunded, and any additional responsibilities charged to the agency should be met with appropriate increases in the resources that are available to the office.

---

6 Cass R. Sunstein, *The Regulatory Lookback*, 94 B.U. L. REV. 579 (2014) (describing author's experience as OIRA administration overseeing lookback process); John D. Graham, *Saving Lives through Administrative Law and Economics*, 157 U. PENN. L. REV. 395, 460 (2008) (discussing prompt letters during author's time as OIRA administrator).

7 *See* Nicholas Bagley & Richard L. Revesz, *Centralized Oversight of the Regulatory State*, 106 COLUM. L. REV. 1260, 1278–79 (2006) (describing deregulatory prompt letters).

8 *See* Alexander Bolton, Rachel Augustine Potter & Sharece Thrower, *Organizational Capacity, Regulatory Review, and the Limits of Political Control*, 32 J. L. ECON. & ORG. 242 (2015) (documenting declining staffing levels and effect of under-staffing on performance).

The remainder of this Essay proceeds as follows. Part I provides some background on the limited role of cost-benefit analysis is agency-agenda setting, discussing the legal, institutional, and political factors that influence how agencies allocate their limited resources. Part II discusses the problem of agency inaction, examining problems of overly active and overly passive agencies. Part III describes why catastrophic risks are particularly prone to generate agency inattention and inaction. This part draws from relevant psychological literature on risk perception, as well as more general political theory concerning agency incentives. Part IV outlines a proposed Catastrophic Risk Review, describing how it can build on successful prior practices, and how it can help integrate cost-benefit analysis into the process of understanding and addressing risks that have the potential to profoundly affect the future of humankind.

I. COST-BENEFIT ANALYSIS AND AGENCY AGENDA SETTING

Among the most important decisions that agencies face is how to allocate their limited capacities. Agencies are charged with a wide range of tasks that include rulemaking and enforcement, permitting, grant oversight, and research. The resources to carry out these tasks, which include personnel, political capital, and financing, are all, of course, scarce. Given the substantial mandates of many agencies, there is a nearly unlimited possible set of priorities that they could adopt. Even in an ideal decision-making environment, determining how to distribute finite resources to best promote public well-being would be a difficult and complex task. And, of course, agencies are not embedded in an ideal decision-making environment. Instead, they must tailor their agendas to fit legal constraints, changing priorities of political overseers, and an often-steady stream of small-scale emergencies.

In practice, there are wide range of influences that affect how agencies devote their time and energies.[9] The most important class of external influences arise from courts, Congress, and the White House. With respect to courts, agencies are continually emmeshed in litigation, which expose them to near-constant oversight by courts. Judges commonly strike down and seek to shape agency actions. The standards set by judicial review—foremost, how courts have interpreted the "arbitrary or capricious" standard under the APA—have created a wide-ranging set of procedural and analytic requirements.[10] These requirements dictate, to a

[9] *See generally* JAMES Q. WILSON, BUREAUCRACY; WHAT GOVERNMENT AGENCIES DO AND WHY THEY DO IT (1989).

[10] There is a considerable literature on the value and burdens associated with probing judicial review of administrative action and the procedural and analytic requirements that come with it. *See* Nicholas Bagley, *The Procedural Fetish*, 118 MICH. L. REV. 345 (2019); Aaron L. Nielson, *Sticky Regulations*, 85 U. CHI. L. REV. 85 (2018); Jason Webb

considerable extent, the resources that agencies must devote to the rulemaking process. The extensive requirements of courts constrain the agenda space of agencies, given a finite rulemaking budget. Courts also review how agencies interpret the laws that they administer and regularly reverse agency actions that are determined to be contrary to statutory authority.

With respect to Congress, the most straightforward mechanism of legislative control is through the law, which determines the scope of agency authority and often prescribes in substantial detail how regulatory activities must take place.[11] Congress also sets agency budgets, and the budget rider process is frequently directed at affecting agency decision making. Agencies must also be responsive to congressional hearings and other oversight activities.

The White House is the third major oversight body. Political scientist Terry Moe has usefully classified the tools that Presidents use to influence agency decision making as either *centralization* or *politicization*.[12] Centralization tools shift the locus of decision making from agencies to the White House. The canonical example of centralization is regulatory review by OIRA, but there is now a substantial presidential bureaucracy that either substitutes for, or at the very least complements, agency decision making. Politicization refers to the increasing trend of Presidents to place loyalists in the most senior managerial positions at agencies. These political appointees are then charged with ensuring that agencies reflect the current priorities of the administration.

Beyond the constitutional branches, there are a wide range of external actors that affect how agencies allocate their scarce resources, in part mediated through courts, Congress, or the White House. Under the APA and other relevant laws, interested parties can petition agencies for rulemakings, and, when those petitions are denied, agencies may be called to defend those decisions in court. A petitioning process under the Endangered Species Act is a common tool used by environmental organizations for directing agency attention to species of concern.[13] In *Massachusetts v. EPA*, the Supreme Court found that the Clean Air Act provided the Environmental Protection Agency (EPA) with the authority to regulate greenhouse gas emissions in

---

Yackee & Susan Webb Yackee, *Testing the Ossification Thesis: An Empirical Examination of Federal Regulatory Volume and Speed*, 80 GEO. WASH. L. REV. 1414 (2012); and Thomas O. McGarity, *Some Thoughts on "Deossifying" the Rulemaking Process*, 41 DUKE L.J. 1385 (1992).

11 Jack M. Beermann, *Congressional Administration*, 43 SAN DIEGO L. REV. 61 (2006).

12 Terry M. Moe, *The Politicized Presidency*, in THE NEW DIRECTION IN AMERICAN POLITICS 235, 235 (John E. Chubb & Paul E. Peterson eds., 1985).

13 Emily E. Puckett, Dylan C. Kesler & D. Noah Greenwald, *Taxa, Petitioning Agency, and Lawsuits Affect Time Spent Awaiting Listing under the US Endangered Species Ac*t, 201 BIOLOGICAL CONSERVATION 220 (2016).

the course of reviewing that agency's denial of a petition for rulemaking.[14] Private parties also decide what agency actions they wish to challenge and whether to pursue political action to attempt to trigger oversight by Congress or the White House.

Agencies also have their own internal reasons to be concerned with the views of interest groups and public opinion more broadly. Good relations between an agency and regulated actors helps facilitate voluntary compliance, an essential in light of limited enforcement budgets.[15] Personnel at agencies may seek to maintain reputations among potential employers to increase their chances of securing lucrative opportunities after their time in government. The staff at agencies may also be concerned about broader public perceptions of the efficacy and usefulness of their organizations.

In addition to external influences, agencies have their own internal processes that help shape their agendas. Some agencies conduct periodic reviews of certain regulatory regimes to update requirements on the basis of new information. New regulatory initiatives are sometimes undertaken as the result of internal study groups or task forces. Agency culture, which tends to remain relatively stable over time, likely also has a role to play in influencing agency agenda setting. Agencies also attend to the work of related government institutions, including other agencies, as well as scientific institutions and assessment bodies, such as those convened by the National Academies. Pressures within the executive is another set of influences on agency agenda setting.[16]

A perhaps striking feature of the various external and internal influences on agency agenda setting is the lack of any formal use of cost-benefit criteria. Some judges may also be more positively disposed toward rules that have overall beneficial effects on society,[17] although the values that individual judges bring to decisions may or may not align with efficiency.[18] It is possible that Congress and the White House have some implicit orientation toward interventions that deliver net benefits, in as

---

14 Massachusetts v. EPA, 549 U.S. 497 (2007).

15 John T. Scholz, *Voluntary Compliance and Regulatory Enforcement*, 6 L. & POLICY 385 (1984).

16 Daphna Renan, *Pooling Powers*, 115 COLUM. L. REV. 211 (2015); J.R. DeShazo & Jody Freeman, *Public Agencies as Lobbyist,* 105 COLUM. L. REV. 2217 (2005).

17 Carolyn Cecot and W. Kip Viscusi, *Judicial Review of Agency Benefit-Cost Analysis*, 22 GEORGE MASON L. REV. 575 (2015); Cass R. Sunstein, *Cost-Benefit Default Principles*, 99 MICH. L. REV. 1651 (2001).

18 For example, libertarian judges may be skeptical of even highly cost-benefit justified government action. Cass R. Sunstein and Adrian Vermeule, *Libertarian Administrative Law*, 82 U. CHI. L. REV. 292 (2015). And, with the rise of textualism, many judges stridently disclaim that policy factors play any role in their process of interpreting legal texts.

much as political officials are judged by voters on the basis of sound government policy. But interest group pressure, partisanship, and symbolic gesturing may be as strong an influence on how political oversight is exercised as any concern with efficiency or welfare maximization. Agencies may have internal norms that favor the prioritization of regulatory actions with substantial welfare payoffs, but there are many internal factors at play beyond regard for social welfare. The case-by-case oversight provided by courts is particularly ill-suited toward pressuring agencies to allocate resources across potential actions in a sensible fashion. There is generally no argument under the APA that an agency's action is arbitrary merely because it reflects poor agenda setting (i.e., that the agency could have engaged in an unrelated action with higher net benefits). These are the kinds of prioritization decisions that courts are inclined to leave to executive discretion.

The upshot of the decision-making environment faced by agencies is that agenda setting is likely to be responsive to various actors, but it is very unlikely to reflect a rational allocation of government resources toward the goal of welfare maximization. Agencies set their agendas to please political overseers, improve their chances of success in court, respond to pressure from interest groups, and remain in relatively good standing in the eyes of the public. Elections, the flux of economic and political events, salient accidents that lead to media attention, and other exigencies often require some form of immediate response, which can distract agencies from longer-term priorities. A particular problem in the contemporary period is the wildly oscillating policy priorities of agencies after changes in the White House.[19] It is not uncommon for a presidential election to result in agency personnel being required to execute an about-face and begin undoing their work over the prior several years.

Reforms to cost-benefit analysis methodology to improve treatment of catastrophic risks can help agencies better understand and manage those risks. But because cost-benefit analysis plays relatively little role placing items on agencies' agendas, those improvements will be mostly relevant for risks that have been identified, characterized, and allocated to an agency through some other process—such as congressional oversight. The processes that affect agency agenda setting are responsive to a wide range of political, economic, and social forces. But there is no formal, and likely little informal, role for cost-benefit analysis in urging agencies to shift from the status quo.

---

[19] Michael A. Livermore & Daniel Richardson, *Administrative Law in an Era of Partisan Volatility*, 69 EMORY L. REV. 1 (2019); Matthew C. Stephenson, *Optimal Political Control of the Bureaucracy*, 10 MICH. L. REV. 53 (2008).

II. THE PROBLEM OF AGENCY INACTION

As discussed in the previous section, the role of cost-benefit analysis on agency decision making is almost entirely as a constraint, not a directive force. OIRA reviews agency cost-benefit analyses and may suggest changes to increase net benefits, but the initiative to start the rulemaking process is with the agency. The alternatives discussed during review are typically relatively minor revisions, not whether a rule should be pursued at all in light of available alternative regulatory emphases. Rationality review by courts similarly focuses on the rule at hand, and courts do not ask agencies to investigate the opportunities of pursing one regulatory matter over another. Agencies may occasionally decline to pursue regulations that they believe are not cost-benefit justified—either based on internal norms or due to the presence of executive and judicial review. But once a rule can be justified on legal and cost-benefit terms, review largely loses its ability to direct agency behavior.

There is an influential line of thought that agencies will tend to be overactive, zealously pursuing regulatory mandates and seeking to maximize their influence whenever possible. If this claim is roughly correct, then concerns about agency prioritization could be dealt with via the use of cost-benefit analysis as a constraint. Agencies would be expected to occupy as much space as they possibly can, and regulatory review would limit that space to cost-benefit-justified rules. The result would be welfare maximizing, in that all cost-benefit-justified regulatory interventions would be identified and exploited. This would hold for interventions to address catastrophic risks as well as other categories of harm. Under the overactivity hypothesis, whenever an existing mandate provides authority to address catastrophic risks, agencies can be expected to pursue regulatory actions in excess of those that would be justified in cost-benefit terms. So long as regulatory review is there to rein in agencies where needed, the system will achieve appropriate levels of risk management.

The problem with the overactive agency hypothesis is that it is more of a caricature than a realistic picture of how agencies work in practice.[20] The structure of agencies and the political environments in which they are embedded creates a complex set of incentives and behaviors. Sometimes, that may lead to overly stringent regulation or harsh enforcement; at other times, agencies will fail to regulate or engage in lax enforcement. Rather than a consistent set of broad tendencies, agency behavior is highly context specific and can depend on a wide range of changing factors, including the political oversight of the day or even the specific inclinations of certain important career personnel.

---

[20] *See generally* Bagley & Revesz, *supra* note 7 at 1282–1304.

Many of the arguments offered in support of the overactivity hypothesis have a foundation in reality, but they also tend to be overly simplistic. For example, economists in the public choice tradition have pointed out that industry actors may seek out government regulation in their sector as a way to erect barriers to entry.[21] And this story likely has some element of truth: it may be that the incumbent players can shape regulatory interventions in ways that limit their exposure to costs while imposing heavy burdens on potential competitors. This is a legitimate concern and may lead to inefficient rules in some contexts. However, although industry may attempt to use its influence to seek out new rules that burden competitors, it may also seek to be free from regulatory costs altogether, shifting the cost of inaction to the broader, less well-organized public. Even accepting the public choice account of industry influence, it is far from clear that the net effect of industry lobbying will be overactive government.

Others have focused on the incentives or preferences of agency personnel to argue in favor of the overactivity hypothesis. Willian Niskanen, for example, argued that agency heads will seek to increase their own salaries and prestige by seeking ever larger budgets and mandates.[22] For Niskanen and others who have taken up this argument, this tendency toward self-aggrandizement is a simple extension of a rational choice model applied to agency heads. But the link between the budget and mandate of an agency, and the utility enjoyed by the agency's leadership, is extremely tenuous.[23] Agencies are not private firms in which pay for senior executives is linked to measures such as share prices or profits. An increase in an agency's budget or the scope of its mandate does not translate in any straightforward way to monetary compensation for senior officials. Whatever benefits that are generated by aggrandizement are more likely to be psychological rather than pecuniary. Seeking out larger budgets and mandates also comes with downsides for agency leadership. The kind of controversy and oversight associated with expanding budgets and mandates would have disutility for many agency heads.

Another possibility related to the Niskanen theory is that agency leadership and staff tend to identify with their agency's mission, which causes them to focus myopically on forwarding that mission at the expense of the broader social good. Again, there is some plausibility to this idea, and this phenomenon may have an influence on agency behavior some of the time. But agency personnel are drawn to

---

21 *See e.g.* Gary S. Becker, *A Theory of Competition Among Pressure Groups for Political Influence*, 98 Q. J. ECON. 371, 372 (1983); Sam Peltzman, *Toward a More General Theory of Regulation*, 19 J.L. & ECON. 211, 212 (1976); George J. Stigler, *The Theory of Economic Regulation*, 2 BELL J. ECON. & MGMT. SCI. 3, 3 (1971).

22 WILLIAM A. NISKANEN, JR., BUREAUCRACY AND REPRESENTATIVE GOVERNMENT (1971).

23 *See generally* Daryl J. Levinson, *Empire-Building Government in Constitutional Law*, 118 HARV. L. REV. 915, 922–23 (2005).

government service for many reasons apart from attraction to an agency's mission. Some may seek out opportunities for public service generally, and not have strong commitments to any particular policy area. Others may be attracted to the pay, benefits, and lifestyle associated with federal government employment and be indifferent to how their work fits into a broader policy program. Some agency heads may even receive utility from undermining the work of the agency and shrinking its budget and mandate—there were several appointees during the Trump administration for whom this appears to have been the case.

Various versions of the overactivity hypothesis also fail to account for the legal and political operating environment of agencies, which often has a strong status quo bias. Agencies face a host of formal and informal barriers that increase the cost and decrease the payoff of major regulatory action. Under the Administrative Procedure Act, agencies have substantial data collection, reason-giving, and public participation requirements. Internal executive requirements concerning interagency cooperation and regulatory review further add time and resource costs to agency action. These analytic and participatory requirements may all be justified, and may lead to improved decision making, but they nonetheless imposed burdens on agency action.

The process of judicial review also adds an extraordinary amount of uncertainty to the regulatory process. Even after years of expensive regulatory process and detailed analysis, even a very well justified rule can be struck down by an ill-disposed appellate court panel. Courts have many options for striking down agency rules, including flaws (or purported flaws) in a rulemaking process or substantive disagreement over the interpretation of a relevant statute. Recent changes in administrative law, such as the expanded "major questions" doctrine announced in *West Virginia v. EPA*, only inject further uncertainty into the process of judicial review.[24] As with the procedural requirements of the APA and the analytic requirements within the executive, judicial review may, overall, be a positive influence on regulatory decision making. But it nonetheless makes rulemaking more difficult and uncertain.

It is worth emphasizing the asymmetrical nature of the legal environment of agency rulemaking. The procedural requirements of the APA are triggered by agency action, but not by agency inaction. Although the APA provides that courts may

[24] That case involved a challenge to an Environmental Protection Agency rule to reduce greenhouse gas emissions from the electricity-generating sector. Under the "major questions doctrine," the Court requires "clear congressional authorization" for agency regulatory authority in certain cases. *West Virginia v. EPA*, 597 U.S. ___, No. 20–1530, slip op. at 19 (U.S. Sup. Ct. June 30, 2022). This sweeping judicial power to strike down agency actions, coupled with the vagueness of the doctrine, creates substantial uncertainty for agencies engaged in rulemakings.

"compel agency action unlawfully withheld or unreasonably delayed,"[25] the burden placed on plaintiffs to justify this remedy is so high that this provision is nearly a dead letter.[26] Even denials of petitions for rulemakings (which are technically an agency action) are given extremely deferential review by agencies.[27] Within the executive, barring the examples discussed in Part IV, agency inaction is not subject to formal centralized oversight.

There is an important set of counterweights that, to some degree, balance the strong status quo bias of agencies' institutional environments. Presidents enter office with policy programs that often rely on agency action. Especially in an era of partisan gridlock in Congress, Presidents must look to executive agencies to pursue their policy goals.[28] There are several offices within the White House charged with coordinating across government agencies to promote the administration's agenda. These entities—such as the Domestic Policy Council and various ad hoc advisors' offices—provide a source of pressure on agencies that encourage them to be active in relevant domains. Presidents are also inclined to act, or be perceived as acting, on the crises of the day, which is another impetus for agency action. And although action can subject an agency to external critique, including by Congress, in certain circumstances, agency inaction will also be criticized. Especially in the case of highly salient negative events—such as the Flint Water Crisis—agency inaction can come under harsh scrutiny.

If these counterweights—which provide agencies with an impetus to action—are sufficiently strong, it is at least possible that in some areas, a system in which cost-benefit analysis acts primarily as a constraint could lead to efficient outcomes. In a policy domain that the White House has prioritized, an agency could face sufficient pressure to act such that the main concern (from a social welfare perspective) is too much, rather than too little, activity. There are good reasons to be skeptical that there is an overall tendency within the administrative state toward overactivity. But whether there are specific policy domains where such a dynamic exists will be a

---

25 5 U.S.C. § 706(1).

26 *See* Heckler v. Chaney, 470 U.S. 821, 832–33 (1985) (finding agency decision not to take enforcement actions unreviewable); Estate of Smith v. Heckler, 747 F.2d 583, 591 (10th Cir. 1984) (stating that § 706(1) only applies in cases where agencies have "ministerial, clearly defined and peremptory" duties).

27 WildEarth Guardians v. EPA, 751 F.3d 649, 651 (2014) (The "standard that governs [] review of an agency's denial of a rulemaking petition" is "extremely limited and highly deferential[.]") (internal quotation marks and citations omitted).

28 *See generally* Elena Kagan, *Presidential Administration*, 114 HARV. L. REV. 2245 (2001), 2272–2319 (discussing presidential efforts to effectuate policy through administration action).

context-specific empirical question. For reasons described in the next Part, catastrophic risk is unlikely to be such a domain.

III. SYSTEMATIC INATTENTION TO CATASTROPHIC RISK

There are three primary reasons why agencies are often likely to underemphasize catastrophic risks. The first is legal. Catastrophic risks are often cross-cutting and are not clearly delegated to specific agencies. The second is political. In a constitutional system that is fundamentally grounded in electoral accountability, long-term, global risks will likely be underprioritized. The third is psychological. Human beings have difficulty reasoning about low-probability events and long time horizons; this leads both voters and government officials to neglect catastrophic risks. A dedicated review process focused on catastrophic risks through a cost-benefit lens, discussed in more detail in Part IV, can help overcome these biases.

Catastrophic harms may come in many shapes, only some of which fit easily in the agency mandates constructed in U.S. law. The response to the covid-19 pandemic is illustrative. Authority over the varied public policy responses to the pandemic was deeply fractured. Vertically, important decisions were vested at all levels, from municipalities through states to the national government. Horizontally, different agencies were charged with different elements of the response, which included outbreak monitoring, public health measures, vaccine and drug development, and deployment of medical resources. Coordination was a constant challenge and pandemic-related policies quickly became embroiled in partisan conflict, which reduced the efficacy of the U.S. response.

The fractured authority over pandemic-related policies not only inhibited efforts to address covid-19 after it emerged. Perhaps more important was the inadequacy of ex-ante planning to address a threat of this kind. The risk of a severe global pandemic was well known within the public health community, and near misses with earlier coronavirus outbreaks had provided ample warning even that this specific family of viruses posed a threat. Nevertheless, the lack of a clear delegation of authority to a competent agency charged with building the necessary capacity and policies to address pandemic risks likely contributed to the chaotic and ineffective U.S. response.

Climate change is an example of a context wherein management of catastrophic risk, even in a near-best-case scenario, the muddled nature of authority and its diffusion between multiple actors raises challenges. Since at least the Clinton administration, the Environmental Protection Agency has interpreted the Clean Air Act to grant it authority to regulate greenhouse gas emissions. That authority was settled in *Massachusetts v. EPA* and the agency has moved forward, in a limited way, with rules to reduce emissions. But national authority over climate policy more generally is more fractured, with the Department of Energy, the Agriculture

Department, the Federal Energy Regulatory Commission, the Nuclear Regulatory Commission, the National Aeronautics and Space Administration, the Federal Emergency Management Agency, and others playing a role in various elements of climate research and policy. Some states—most notably California—have played a leading role as well, in part because the federal government's response has been weak and inconsistent. And even EPA's authority is under a cloud of uncertainty after the Supreme Court's recent decision in *West Virginia*.[29]

For other catastrophic risks, the situation is even more dire. There is no agency like the EPA that has taken a leading role in addressing advanced artificial intelligence, bioengineering, or other sources of catastrophic risk. In many instances, the concern is not that authority is too diffused—as is likely the case for pandemic risks. Rather, authority is sufficiently unclear that it is difficult to even identify a default actor whose responsibility it would be to recognize and begin to consider the relevant risk. This missing responsibility makes inattention to catastrophic risks all the more likely.

The nature of the U.S. political process also biases government decision makers toward inattention to many catastrophic risks. Large-scale nuclear war, climate change, artificial intelligence, severe global pandemics, runaway bioengineering, and asteroid strikes all present substantial scale mismatches with national electoral democracy. These scale mismatches arise because the vast majority of stakeholders who are affected by decisions concerning these risks fall outside the voting public.

The first scale mismatch is spatial. Many, if not all, catastrophic risks are global in nature. To qualify as catastrophic, a risk must place an extremely large number of people in harm's way and threaten the continued viability of human development, modern civilization, or humankind. For national-level political processes, most of the harms caused by actions that exacerbate catastrophic risks, or benefits that result from steps to mitigate these risks, will fall extra-jurisdictionally. We can expect rational self-interested politicians and voters to underinvest attention in these issues, from the perspective of global welfare. More salient, localized issues, such as employment, economic growth, crime, and education, will likely continue to win out for attention compared to catastrophic risks whose effects will occur on a global scale.

The covid-19 pandemic and climate change again illustrate the natural domestic myopia of national institutions. Despite the clearly global nature of these threats, policies made at the national level have tended to favor local interests and disregard cross-border effects. With respect to the pandemic, vaccine supplies were, and continue to be, distributed extremely unevenly across the globe, with the richer countries experiencing an oversupply, and developing countries continuing to struggle to gain access to high-quality vaccines. This situation has created opportunities for the virus to mutate, leading to successive rounds of variants that

[29] *See supra* note 24.

have increasing ability to evade existing vaccines. A global approach would have focused vaccine supply in a more efficient manner to maximize total immunity and reduce the risks caused by new variants.

With respect to climate change, the debate in the United States over the social cost of carbon encapsulates the difficulties associated with national-level policymakers addressing global catastrophic challenges. The social cost of carbon is a monetary estimate of the damages associated with greenhouse gas emissions. There are a number of contested questions that arise when estimating the social cost of carbon, including how best to model damages and the discount rate to use for harms in the future. In the United States, the question of whether the social cost of carbon should include all damages, or only those that occur domestically, has also arisen. The Obama and Biden administrations have used a global estimate, while the Trump administration favored a domestic-only social cost of carbon. Without debating the economic, legal, or ethical merits of the alternative approaches to the social cost of carbon, the controversy demonstrates that at least some of the time, domestic political institutions can be expected to discount, or even disregard entirely, the global consequences of their actions.[30]

The second scale mismatch is temporal. Most catastrophic risks are long term and relatively unlikely to be realized within the next few decades. But even very small risks, if they are consistent, are translated into near certainties over time. For example, if each year, humanity faces a one-in-a-million risk of an asteroid strike, then the risk of occurrence within a timespan that is meaningful for politics, which operates at a decadal basis at most, would still be very small. But over a 50,000-year time horizon (roughly the period of time since the Cognitive Revolution) the same risk translates into, approximately, a 5% probability of a strike at some point. Political cycles are entirely out of sync with these time spans, and seemingly rational treatment of *de minimis* probabilities results in the neglect of very meaningful risks over the long term.

If a catastrophic risk were realized, it would also impose costs indefinitely. An artificial super-intelligence that was indifferent to human well-being might convert massive amounts of resources to its own aims, substantially reducing the prospects of human development.[31] Such an occurrence would likely be irreversible, as the super-intelligence would continue to advance while humankind stagnated. The prospects of generation after generation of people would be severely diminished, potential affecting many billions of persons. But nearly all (or all) of the persons at risk from these harms have not yet been born, and they are certainly not voting

---

30 *See* Arden Rowell, *Foreign Impacts and Climate Change*, 39 HARV. ENVTL. L. REV. 371 (2015) (arguing that some statutes may require that agencies only consider domestic effects).

31 *See generally* NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES (2014).

members of the public. The normal functioning of democratic institutions will orient them toward the concerns of existing voters, not such future persons.

Finally, the human beings that make up the voting public, politicians, and agency personnel are all subject to psychological tendencies that incline them to underemphasize certain types of catastrophic risks. There is a substantial body of research that documents the difficulties that people have with thinking in probabilistic terms, especially when the relevant probabilities are very small.[32] People will sometimes ignore events whose probabilities fall below a certain threshold.[33] In the case of many catastrophic risks, the probability that any would occur in a given person's lifespan is quite small, potentially leading to the neglect of this risk altogether. The consequences of catastrophic risks—such as nuclear war, large-scale environmental destruction, or asteroid strike—may be so grave that people may avoid thinking about them in order to prevent negative emotional experiences.[34] These risks can also be difficult to manage, leading to feelings of ineffectiveness and helplessness—both negative emotions that people will tend to avoid if possible.[35] The solutions that do exist may also involve institutional or policy choices that people may find unpleasant.[36] Increased government supervision of technological development—which may be necessary to monitor and mitigate the risks associated with artificial intelligence—may strike many people as intrusive and contrary to values of economic liberty. Global cooperation to address climate change may challenge people's national-level identify affiliation. Aversion to such solutions may lead people to downplay or deemphasize the risk in question, so that they might avoid the cognitive dissonance associated with making tragic choices between undertaking necessary but disagreeable steps to address a catastrophic risk or making an informed and clear-eyed choice not to.

It bears mentioning that the overall effect of many of these psychological tendencies is ambiguous, and some of them may result in people paying too much

---

[32] *See* Cass R. Sunstein, *Probability Neglect: Emotions, Worst Cases, and Law*, 112 YALE L. J. 61 (2002).

[33] Colin F. Camerer & Howard Kunreuther, *Decision Processes for Low Probability Events: Policy Implications*, 8 J. POL. ANALYSIS & MGMT. 565, 570 (1989).

[34] *See* Saffron O'Neill & Sophie Nicholson-Cole, *"Fear Won't Do It"" Promoting Positive Engagement with Climate Change through Visual and Iconic Representations*, 30 SCI. COMM. 355 (2009); Robert A.C. Ruiter et al, *Scary Warnings and Rational Precautions: A Review of the Psychology of Fear Appeals*, 66 PSYCH. & HEALTH 613 (2001).

[35] Erika Salomon et al., *Climate Change Helplessness and the (De)moralization of Individual Energy Behavior* 23 J. EXP. PSYCH.: APPLIED 15 (2017).

[36] *See generally* ARDEN ROWELL & KENWORTHEY BILZ, THE PSYCHOLOGY OF ENVIRONMENTAL LAW 230 (2021) (discussing literature on solution aversion in the context of climate change).

attention to certain types of catastrophic risks. For example, a substantial body of research shows that people are inclined to overestimate risks when they are salient.[37] The entertainment industry plays a role in vividly rendering some catastrophic risks in realistic terms that are easily accessible to the human imagination.[38] Dreaded risks also tend to be overestimated.[39] But at least some of the time, human psychology will lead relevant decisionmakers to avoid focusing on certain catastrophic risks, even when it would be rational to do so.

IV. CATASTROPHIC RISK REVIEW

The prior three sections have established the problem. Currently, there is little formal role for cost-benefit analysis in agency agenda setting. Catastrophic risks, like all other risks that are relevant for government decision making, are primarily weighed in a cost-benefit fashion after agency agendas have already been firmly set. This state of affairs, in which the role of cost-benefit analysis is to constrain agency behavior, would be acceptable if agencies could be expected to be generally overactive, expanding and pushing their authority to the maximum extent possible. But this is the not the case, especially so for catastrophic risks. Agencies are complex entities operating in complex political environments. Some of the time, agencies may tend toward overactivity, but there is at least as much reason to be concerned with agency inaction as with inefficient action. And catastrophic risks have several legal, political, and psychological dimensions that make inattention particularly likely.

The proposal in this Essay is to address this problem through a new executive review process focused on catastrophic risks. OIRA would coordinate an interagency process in which catastrophic risks could be identified and policy alternatives could be evaluated using cost-benefit criteria. This new process would be modeled on prior experiences with OIRA and agency agenda setting. Catastrophic Risk Review could also provide the occasion for the development and reform of cost-benefit analysis methodologies that were tailored to that particular context.

Although cost-benefit analysis and OIRA review has generally been reactive—relying on other processes to set agency agendas and then using cost-benefit analysis to steer agencies on courses that are already set—there are some notable exceptions. Administrations of both political parties have employed OIRA in innovative ways to

---

37 Shawn J. McCoy & Randall P. Walsh, *Wildfire Risk, Salience & Housing Demand*, 91 J. ENVTL. ECON. & MGMT. 203 (2018).

38 For example, popular science fiction tales of dinosaurs run amok may lead to some people to overestimate the risks associated with de-extinction efforts. *See e.g.* MICHAEL CRICHTON, JURASSIC PARK (1990).

39 Paul Slovic & Ellen Peters, *Risk Perception and Affect*, 15 CURR. DIR. PSYCH. SCI. 322 (2006).

help prod agencies to action. This more proactive posture for OIRA has met with notable success, and these experiences can help establish a template for a Catastrophic Risk Review process.

An important example of OIRA taking on a more proactive role was the practice of issuing "prompt letters" under the administration of George W. Bush. Carried out during the term of OIRA Administrator John Graham, prompt letters were public requests by OIRA to an agency for some kind of agency action.[40] During the Bush administration a dozen prompt letters were issued on a wide range of issues, including one urging the National Highway Traffic Safety Administration to examine rules to reduce harms from off-center automobile collisions, and another requesting that the Occupational Safety and Health Administration investigate the benefits of requiring automated external defibrillators (AEDs) in the workplace. Perhaps the most successful prompt letter was a request to the Food and Drug Administration to move forward with a labeling requirement for trans-fat content in food.[41] Prompt letters were not always used to promote additional regulation; one important letter issued by Graham encouraged agencies to review a set of several dozen rules that, OIRA argued, imposed inefficient costs.[42]

OIRA undertook an even more substantial and sustained proactive role during the Obama administration. Midway through his first term in office, and with economic effects of the Great Recession still lingering, President Obama issued an executive order requiring agencies to initiate a process of retrospective analysis to identify rules to be "modified, streamlined, expanded, or repealed."[43] OIRA issued guidance on how to conduct this regulatory lookback, and agencies were given deadlines to submit preliminary plans, collect public feedback, and issue final plans to update, change, or rescind rules.[44] Eventually, agencies identified several hundred initiatives that promised billions of dollars in net savings.[45] This regulatory lookback built on efforts of earlier presidents to coordinate similar government-wide reassessment of the stock of existing rules.[46]

---

40 *See* Graham, *supra* note 6 at 460.

41 Graham credits this prompt letter with initiating a broader set of social policies aimed at reducing trans-fat content. *See id.*

42 *See* Bagley & Revesz *supra* note 7.

43 Exec. Order 13,563, 76 Fed. Reg 3821 (January 18, 2011).

44 Memorandum from Cass R. Sunstein, Adm'r, Office of Info. & Regulatory Affairs, to the Heads of Exec. Dep'ts & Agencies (Apr. 25, 2011), archived at http://perma.cc/X9MQ-K88W.

45 Sunstein *supra* note 6 at 593.

46 Cary Coglianese, *Moving Forward with Regulatory Lookback*, 30 YALE J. REG. ONLINE 57, 58–59 (2012).

Although neither the prompt letters nor the regulatory lookback generated overwhelming outcomes, they were nonetheless useful innovations that had important successes. It is likely that important rules were adopted, or at least expedited, through the prompt letter process. Similarly, the regulatory lookback redirected agency resources to a reassessment process that identified reforms that likely would have languished. Of course, it is difficult to determine, overall, whether these efforts were wise uses of agency resources. It is possible that by diverting agency attention from other work to respond to these central demands, other even more net beneficial projects were delayed. Indeed, both the prompt letters and the regulatory lookback were responsive in part to the political exigencies of the day: during the Bush administration, to not appear overly anti-regulatory, and during the Obama administration, to push back against claims that agencies were prone to overregulate. Regardless of their overall value, however, prompt letters and regulatory lookback provided a proof of concept that OIRA can use its position in a proactive fashion, in line with cost-benefit principles, when an administration is interested in doing so.

Other processes that provide some precedent for an OIRA-led Catastrophic Risk Review are government-wide guidance documents that the Office has issued. OIRA was the lead coordinator of the interagency process that developed the social cost of carbon, used by agencies across the federal government when estimating the effects of decisions with consequences for greenhouse gas emissions. OIRA also issued Circular A-4, which establishes best practices for conducting cost-benefit analysis across the government. In addition, OIRA is charged with issuing annual reports to Congress on the costs and benefits of federal rulemaking. These reports have historically provided the office with opportunities to highlight important open questions in cost-benefit analysis methodology and risk regulation.

Based on these prior experiences, there are some important features that Catastrophic Risk Review should have. Ideally, as was the case for the regulatory lookback, it should be initiated by the President in an executive order. Presidential leadership can help establish the importance of the process and ensure adequate agency attention. It should also be government-wide and systematic, again similar to the regulatory lookback and social cost of carbon processes. Although the prompt letters were valuable, they were also ad hoc. It was easier to dismiss prompt letters as an artifact of the personal views of a particular OIRA administration, rather than as reflecting a widely shared, cross-administration consensus. OIRA can play a coordinating and oversight role, and OIRA personnel could staff the project, but an interagency coordinating group should also be appointed to elicit expertise and generate buy-in across the executive.

One feature of the prompt letters that Catastrophic Risk Review should adopt is a focus on identifying high net present value undertakings. Although OIRA has an institutional orientation toward net-benefit maximization, there have been instances

when the agency was directed to ignore net benefits in favor of other priorities. For example, Trump administration executive orders required agencies to impose no new net costs through regulation, and to rescind two rules for every one rule that they issued. Neither one of these requirements included any language to ensure that they were implemented in such a way as to maximize net benefits. These Trump-era requirements have been criticized as irrational and inconsistent with OIRA's larger mission.[47] Were Catastrophic Risk Review oriented toward the precautionary principle or some other criteria that was not welfare maximization, there is some risk that it could be viewed similarly as outside OIRA's purview and conflicting with deep normative principles embedded in the U.S. administrative state.

Prior experience indicated that there should also be a central role for public participation in Catastrophic Risk Review. Two successful cases of integrating public participation into OIRA-led cross-government efforts were in the process of developing Circular A-4, and a public solicitation during the early Obama administration on improving the process of regulatory review. Circular A-4 developed a draft document that was subject to interagency review, public comment, and peer review. This process helped improve and legitimate the document, and the long influence of Circular A-4 helps demonstrate the value of this open and public process. The Obama administration also issued a public call for comments on the process of regulatory review generally, and it received a substantial number of suggestions from a diverse set of commenters. Unfortunately, that process did not ultimately produce substantive reforms, but it nonetheless created a forum to collect innovative ideas from interested parties. By contrast to these two more public processes, the interagency working group on the social cost of carbon was an internal government process that did not subject its work to public comment or peer review. The social cost of carbon has met with some criticism due to the non-public nature of its development. It is likely the case that, at the very least from a legitimacy perspective, the social cost of carbon adopted during the Obama administration would have benefited from opportunities for public comment and peer review.[48]

Given the cross-cutting nature of catastrophic risks, and the reality that many of these issues do not map cleanly onto existing agency mandates, it would be useful for Catastrophic Risk Review to be explicitly oriented to identify a range of

---

47 *See* Jodi Short, *The Trouble with Counting: Cutting through the Rhetoric of Red Tape Cutting*, 103 MINN. L. REV. 93, 94 n.3 (2018) (collecting criticisms).

48 It bears noting that any time the social cost of carbon is used in a regulatory process, the requirements of the Administrative Procedure Act ensure that the public has an opportunity to comment. Nevertheless, the social cost of carbon estimates themselves were developed through an internal process that was criticized as being overly secretive. Opportunities for comment would have cut off these concerns and may have generated useful information for the working group.

government policy options, with potentially different institutional audiences. Prompt letters and regulatory lookback sought to identify agency action—regulatory or deregulatory—and a list of actions for agencies to consider would certainly be a useful output from Catastrophic Risk Review. But the review process could also identify policies that require congressional action, and these findings could be summarized in a document such as OIRA's annual report to Congress. Catastrophic Risk Review could also identify areas where funding for research through entities like the National Science Foundation would be appropriate, or where assessment reports by bodies such as the National Academy of Sciences could usefully inform policy. Catastrophic Risk Review could also identify areas of cost-benefit analysis methodology that should be updated in light of catastrophic risks.

After an initial period of heightened activity, Catastrophic Risk Review should be incorporated into a long-term, ongoing process of identifying and addressing catastrophic risks. Given the historical inattention to catastrophic risks, there is a substantial backlog of unaddressed issues. The goal of the first stage of Catastrophic Risk Review should be to address this backlog, identify outstanding issues, and offer recommendations for additional measures. Once this initial step has taken place, Catastrophic Risk Review should transition into a more regularized process of information gathering and agency coordination, perhaps punctuated with periodic review processes with more intensive staffing and public participation.

To summarize the process just described, Catastrophic Risk Review would be a government-wide and systematic review of a range of catastrophic risks, including environmental, technological, and socio-economic risks, that was initiated at the presidential level and supervised by OIRA. The review process would include a substantial public participation component and would be charged with examining and assessing a range of policy responses from a cost-benefit perspective and making recommendations to agencies, Congress, and other relevant bodies. Such a process would be a major innovation in how the United States addresses catastrophic risks and would move these issues from the periphery of government decision making toward a more central place in setting the long-term agenda for the U.S. policy process.

As should be clear, such a process would require considerable resources. It bears emphasis that OIRA is already severely stretched in its capacity to address the day-to-day work of the Office. The task of reviewing regulatory proposals from all of the major agencies is extremely demanding. Not only is technical expertise required, but considerable time must be spent coordinating across the White House and executive branch. Simply adding an additional task to OIRA's already demanding mandate is unlikely to result in a useful and thorough process. Any additional work burden that is placed on OIRA should be accompanied by appropriate resources to ensure that the task can be carried out in a timely and careful manner.

CONCLUSION

The greatest impediment to the use of cost-benefit analysis to improve the management of catastrophic risks is the reactive nature of regulatory review. In the current regulatory environment, agencies are relatively unlikely to exacerbate catastrophic risks through their actions. Rather, the greatest contribution of the regulatory system to catastrophic harms is through inattention and inaction. Improving cost-benefit analysis methodology, without accompanying procedural reforms, will not remedy this fundamental problem.

This Essay has proposed a new Catastrophic Risk Review process that would focus OIRA and administrative agencies on the task of identifying catastrophic risks and evaluating steps that could be taken to address them. If possible, such a process should be initiated at the highest levels, should be government-wide in scope, should involve substantial public participation, and should contemplate a wide range of risks and potential policy responses. This Catastrophic Risk Review builds on earlier successes at OIRA in prompting regulatory action and coordination activity across the executive.

Although many catastrophic risks have relatively small probabilities, the scale of their harms, were they to come to fruition, are enormous. Human psychology and institutions are ill-suited to recognizing and managing such risks. But methods like cost-benefit analysis, and processes like regulatory review, persist exactly because they complement and mitigate these cognitive and institutional limitations. Cost-benefit analysis can help clarify and debias our understanding of the scale and consequences of catastrophic risks, and the value—and costs—of efforts to manage these risks. But to be useful, even the best cost-benefit analysis methodology must be embedded in a process in which it is carried out and its insights can inform meaningful policy. Catastrophic Risk Review would provide a forum in which the tool of cost-benefit analysis can be put to good use to identify and support sound government policies to address harms that threaten the continued viability of the human project.