Why Are Pandemics Ideological?

Tom S. Clark John W. Patty

July 31, 2020

Abstract

Effective governmental responses to disasters rely in part on the expertise and skills of government workers. Building and retaining expertise within the government often requires granting unelected civil servants discretion over policy-relevant decisions. We present a simple model of policy-making that captures the trade-offs faced by a policy-maker when considering implementing a policy that may reduce the level of expertise within the government in the shadow of a potential disaster in the future. The model illustrates how policy motivations of an elected government might lead to governmental failure in a response to a future disaster. The model predicts that governmental failures will be more likely when the policy-maker has relatively extreme preferences, and that the failures will tend to occur in agencies where the experts' policy preferences are opposed to those of the policy-maker.

1 Introduction

Public catastrophes are often marked by periods of political unity. In the aftermath of the September 11 terrorist attacks, Americans from disparate political orientations came together to support the government's response, and the Great Recession saw a broad collaborative effort to respond to economic tragedy. Indeed, an empirical finding from the extant literature is that American presidents benefit from the lowering of partisan barriers in periods of national strife. The global health pandemic brought on by the spread of a novel coronavirus in 2020 ought to be no exception. It has prompted large-scale, mass coordination efforts within and among countries around the world. However, as distinct from militarized threats to national security or more typical periods of economic pain, the pandemic has devolved in the United States into an all-too-familiar pattern of hyper-partisan debate. Why might a pandemic become politicized, and why might leaders put politics ahead of public health and the national good?

Over at least the last half century, Americans have witnessed "a revival of Jacksonian hostility toward expertise, and of the belief that common sense is an adequate substitute for technical knowledge" (Nelkin, 1975, 36). This observation was particularly keen in the 1970s as the public, policy-makers, and scientists debated a plan for American energy infrastructure amidst concern about environmental safety. However, the pattern of politicizing scientific knowledge was neither new nor at its climax in the 1970s. In the 21st century, Americans are witnessing just as much—if not more—political hand-wringing over the value of expertise in policy-making. By the beginning of the century, commentators and critics were noting a troubling and seemingly systematic rejection of intellectualism and rationality (e.g., Jacoby, 2018). While it is unclear who leads whom in this path towards polarized factual foundations, the anti-scientific movement has affected political issues ranging from evolution to climate change to foreign policy to public health. Indeed, from the promise of Donald Trump to "drain the swamp" to prominent leaders of the "anti-vaccine" movement to conspiracy groups, such as QAnon, American politics today is very much shaped by the leadership of cavalier-outsiderism that prioritizes a rejection of expertise and priority on skepticism.[1]

But this general trend does not answer the question why pandemics should be so ideological. Is there something special about moments of extreme urgency that exacerbate the politicization of science precisely when we ought to be activated by shared values and a common goal (e.g., saving life and overcoming a pandemic)? Given there ought not to be competing agendas, shouldn't a global health pandemic be precisely the setting where we should expect politicians to lower the barriers to implementing scientific knowledge? We are not. Instead, we argue that it is precisely in these settings where we ought to see heightened partisan divides driven by structural incentives that polarized leaders have to undermine the scientists who might otherwise help a united effort to overcome tragedy.

We argue that in order to retain (and attract) high-quality experts who can provide valuable technical and scientific capacity in the event of a disaster, and such policy concessions can be too much for some politicians to bear. In particular, ideological extremists will find it costlier to moderate policy enough to retain a given expert, relative to more moderate politicians. That incentive interacts with the ideological bias of political leaders to differentially affect the scientific expertise with which administrative agencies are endowed. We will therefore expect disasters under ideological extremists to be less likely to benefit from expertise that can mitigate the extent of the disaster. Crucially, our model is consistent with two different avenues of risk mitigation—preventing disasters or minimizing the destruction caused by a disaster once it happens. The result is that disasters are more likely to happen under extremists, and conditional upon a disaster occurring, it will be more disastrous under extremists leaders than under more moderate leaders.

In our view, the function of expertise in the context of a disaster is that it can help provide the capacity to effectively respond and mitigate the damage caused by the event.

[1]Indeed, in many ways, Stephen Colbert captured well this characteristic of American politics in 2005 during the premier of his show, *The Colbert Report*, in which he defined "truthiness" as a characteristic of a statement's truth value not being in relation to its veracity but instead in relation to what one believes.

For example, in the context of the covid-19 pandemic, once the disaster occurred—that is, once the novel coronavirus began spreading among people—scientific expertise in government facilitated the identification of effective measures to contain its spread and adopt policies that would limit its health and economic ramifications. The model we develop captures this dynamic and examines the implications of a political tradeoff between (a) offering policy concessions *ex ante* in order to have access to scientific expertise in the event of a disaster and (b) picking a less-constrained policy but foregoing bureaucratic expertise.

The article is organized as follows. In the next section, we describe a series of findings from the extant literature that provide the intellectual foundations for our analysis. We then present a formalization of a baseline model that isolates the strategic tensions policy-makers face in the context of scientific expertise and mitigation of risk from disasters. We evaluate equilibrium predictions from the model and derive a series of empirical implications. We then investigate how patterns in outcomes during disasters, such as pandemics, are affected by the ideological orientation of policy-makers, with a particular focus on extremists, as opposed to moderate, political leadership. We then discuss a series of extensions to the baseline model and conclude with discussion of how our model informs debates about policy-making and technical and bureaucratic capacity.

2 Expertise and Policy-Making

Scholars of policy-making and bureaucratic capacity have long been interested in the conditions under which politicians are able to benefit from and harvest technical and scientific expertise to enhance the functioning of government. Indeed, the role of academics, scientists, and experts has been an issue at the heart of a core American tension—the value of academic, as opposed to practical, knowledge for governing (e.g., Hofstadter, 1963). We build from two key insights. First, scientific capacity does not exogenously emerge within an institution, and policy-makers must make choices that create incentives for bureaucrats to develop and retain scientific expertise (e.g., Perry and Wise, 1990; Dewatripont, Jewitt and Tirole, 1999; Gailmard and Patty, 2007a; Cameron and de Figueiredo, 2018; Gailmard and Patty, 2012). Second, Politicians and administrative agencies typically have ideological biases (e.g., Bertelli and Grose, 2009; Richardson, Clinton and Lewis, 2018; Clinton and Lewis, 2008; Snyder and Weingast, 2000). These ideological biases mean that across different policy domains, a political leader may find herself more or less aligned with an agency's mission. We believe that feature of governance can interact with the need to foster bureaucratic competence in the first place. The interaction of those two dynamics is at the core of our analysis. As we illustrate, together they suggest an answer to why we might observe ideological polarization in the context of disasters, such as pandemics.

In brief, the function of expertise in the context of a disaster is that it can help provide the capacity to effectively respond and mitigate the damage caused by the event. In the context of the covid-19 pandemic, once the disaster occurred—that is, once the novel coronavirus began spreading among people—scientific expertise in government facilitated the identification of effective measures to contain its spread and adopt policies that would limit its health and economic ramifications. The model we develop captures this dynamic and examines the implications of a political tradeoff between (a) offering policy concessions *ex ante* in order to have access to scientific expertise in the event of a disaster and (b) picking a less-constrained policy but foregoing bureaucratic expertise.

We now turn to a formalization of the politics of expertise and disasters. We build our model by first isolating the core tension we derive from the literature. We then discuss a couple of extensions to evaluate how institutional features, such as collegial policy making and audits of agency effectiveness affect the incentives we find in the baseline model. We show that throughout the various institutional settings two relationships emerge. First, conditional upon a disaster happening, it is more likely to be "bad" when an ideological extremist is in charge of policy than when an ideological moderate is. Second, ideological extremists are more likely to have disasters happen than are ideological moderates.

2.1 Scientific Expertise and Constraints on Policy-Making

Scholars have extensively studied the role that expertise plays in policy-making. Typically, a lot of analytic focus is centered on how bureaucratic expertise can facilitate a politician's ability to make policy that will achieve her objective. One core lesson from that literature is known as the ally principle, which states that "as the policy preferences of politicians and bureaucrats converge, politicians will delegate more discretion to bureaucrats" (Huber and Shipan, 2006, 853). The key mechanism underlying this result is that because with discretion bureaucrats can cause policy outcomes to move towards their preferred policy, it is only when politicians and bureaucrats agree that politicians will want bureaucrats to be able to use their discretion to influence the implementation of policy.

A related lesson is that expertise is neither necessarily available nor free to access. Politicians must incentivize a bureaucracy to develop expertise, and extracting information from potential technical experts usually entails conceding some discretion over policy (e.g., Gailmard and Patty, 2012). One way in which that constraint manifests is through delegation, by which well-understood problems of agency loss can easily arise. However, a second way in which technical, scientific expertise can constrain a policy-maker is through labor market forces. Scientific knowledge is valuable, and as a consequence, bureaucratic experts—especially the ones with valuable information—typically have outside career options. They also typically care about their subject matter.[2]

Consider an example. In January 2017, Donald Trump became president and replaced then-Acting Administrator Catherine McCabe with Scott Pruitt, a widely-known opponent of environmental protection regulations. Pruitt induced a conservative shift in the EPA's policies and mission. By the end of the year, national media were reporting on both major changes in the EPA's ideological orientation[3] as well as massive departures from the ranks

[2]In innate interest in a subject matter might be the very reason they choose their field of expertise in the first place.

[3]Dennis, Brady; Eilperin, Juliet (December 31, 2017). "How Scott Pruitt turned the EPA into one of Trump's most powerful tools". Washington Post.

of career scientists and officials at the EPA.[4] By late 2019, the EPA was stating that the water in Flint was safe to drink, while scientists outside the EPA were casting doubt on that claim.[5] The in ability to conclusively resolve the Flint water crisis more than 4 years after it began to attract the attention of a national bureaucracy is, in our argument, in part a function of the fact that the conservative shift in the EPA's policies and mission under the Trump Administration induced scientists and related experts to leave the agency.

Beyond the EPA example, there is ample reason to believe that the adoption of extreme policies—especially ones inconsistent with the perceived mission of a bureaucracy—induces technical and scientific experts to depart from government service. A wide range of theoretical models have examined the incentives for bureaucrats to depart when there is conflict over policy objectives with political principals (e.g., Perry and Wise, 1990; Dewatripont, Jewitt and Tirole, 1999; Gailmard and Patty, 2007*a*; Cameron and de Figueiredo, 2018). These models typically illustrate that the development or acquisition of policy-relevant expertise is linked to the way in which policy-makers will implement that information, creating a trade-off between capacity and control on the part of political principals, as in the EPA example, What is more, empirical evidence suggests that the relative attractiveness of outside options and the political conditions of the job influence civil servants' decisions to depart from public service (e.g., Bertelli and Lewis, 2012; Doherty, Lewis and Limbocker, 2019; Bolton, De Figueiredo and Lewis, 2016).

As a consequence, in order to retain (and attract) high-quality technical experts, policy-makers might find that there is a limit to how far they can deviate from the kinds of policies the expert finds acceptable. The recent experiences of life-time civil servants, ranging from bureaucrats in the Departments of Justice and State to the Environmental Protection Agency and Centers for Disease Control and Prevention—illustrates what can happen when political policy-makers push policy and practice to a point where technical experts decide to "throw their hands up" and leave government.

[4]https://www.nytimes.com/2017/12/22/climate/epa-buyouts-pruitt.html

[5]https://www.pbs.org/wgbh/frontline/article/epa-says-flints-water-is-safe-scientists-arent-so-sure/

The effect of this dynamic is that there can be situations in which the value of expertise is not justified by the policy concessions necessary to realize it. But, how does the rejection of expertise affect disasters? In addition to the global health pandemic associated with covid-19, we consider two examples to help illustrate the underlying mechanisms. In 2016, the Environmental Protection Agency issued a "Safe Drinking Water Act Emergency Order" in Flint, Michigan, advising residents that their tap water was unhealthy. Over the course of that year, several steps were taken to identify the causes and potential solutions to the drinking water problems in Flint. A major component of those efforts involved assessing whether improvements to the drinking water infrastructure were sufficient to render it safe.[6]

2.2 Ideological Bias and Scientific Expertise

To the extent we are focused on the decision of expert bureaucrats to remain in government service, it is useful to consider why we might find variation in ideological conflict between policy-makers and technical experts. In our view, the incentives created by policy choices for scientific experts' career decisions are potentially exacerbated by ideological bias by policy-makers. An extensive literature has examined the tension policy-makers confront between employing experts and employing "loyalists" (e.g., Moe, 1985; Rudalevige and Lewis, 2005; Rudalevige, 2009; Krause and O'Connell, 2019; Carpenter and Krause, 2014; Lewis, 2009). That literature focuses on a tension that leaders face between endowing an agency—more generically, any agent—with independent capacity or, alternatively, tightening control at the expense of flexibility and capacity. A related, but conceptually distinct, literature examines the extent to which political leaders' ideological biases influence how they direct government spending (e.g., Gordon, 2011). As with the tension between loyalty and expertise, a core insight from these studies is that ideological biases shape the way in which policy is carried

[6] Of course, the role of scientific expertise in disasters is not limited to policy responses; one might imagine the scientific expertise can contribute to the risk of a disaster happening in the first place. The model we develop below can account for both dynamics, because we assume that politicians view scientific expertise as risk mitigation. We model the role of expertise as decreasing the damage in the event a disaster occurs, but the functional form we select is equivalent to one in which expertise reduces the risk of a disaster in the first instance.

out, potentially creating inefficiencies or sub-optimal governance. One goal of our analysis is to model the interaction of these alternative incentives to assess whether they are in tension with each other and how they affect empirical expectations about policy-maker behavior and the politics of disasters.

Of particular interest in our model is how conditions of ideological polarization interact with the politics of bureaucratic expertise these literatures study. One finding that helps provide traction, from which we build is that the various missions of administrative agencies interact with political ideology differently. A variety of empirical studies have documented variation in perceptions of the underlying ideological orientation of various administrative agencies (Bertelli and Grose, 2009; Richardson, Clinton and Lewis, 2018; Clinton and Lewis, 2008; Snyder and Weingast, 2000). The implication, then, is that the extent to which the policy-making-bureaucratic-career tension is activated and shapes policy might be influenced by ideological bias among policy-makers. Therefore, we might expect that political turnover that shifts an agency's mission or ideological orientation might result in a loss of bureaucratic expertise or capacity. That dynamic should be particularly acute when turnover is associated with the installation of an ideologically extreme principal who advances objectives that are inconsistent with the preferences of most scientific experts. In the event, then, of a disaster, such as a global health pandemic, ideologically extreme leaders might have less technical capacity to respond to the emergency. Is is this interaction that drives the core tension we model below.

2.3 Politics, Disasters, and Risk Mitigation

These twin lessons from the study of bureaucracy in the policy-making process suggest an interesting puzzle, because they reveal that expertise and information are not unambiguously good for policy-makers. We build from these insights to address a simple question: why might we observe partisan polarization in the face of national emergencies that seemingly create opportunities for Pareto improvements for everyone? In contrast to periods of more "normal" politics, global health pandemics and other national emergencies ought to unify citizens, as we have seen in the wake of major terrorist attacks or natural disasters. It is our assumption that in the wake of such a disaster, everyone would be made better off with some kind of policy change. There might not be agreement about which policy change is better, but we assume that, for example, everyone would prefer a world without an active pandemic to a world with an active pandemic. What explains divergent patterns of partisanship in the face of national disasters? What role does ideological and bureaucratic competence play in the consequences of national emergencies?

To answer the question, we must move beyond standard models of bureaucratic expertise in the shadow of ideological policy-making. In particular, by building from models of bureaucratic expertise and capacity, we show that the nature of ideological extremity means there is a wider range of scientists who an extremist would rather "lose" than make policy concessions to retain than would an ideological moderate. As a result, we should expect ideological extremists, *ceteris paribus*, to be more likely to adopt policies that induce policy experts to depart from government service, undermining their capacity to respond to disasters, should they occur. Of course, the world of bureaucratic expertise and policy-making is more complicated than a single principal making policy in the shadow of bureaucratic expertise. Most obviously, legislatures play a particularly important role in policy-making (e.g., Huber, Shipan and Pfahler, 2001; Huber and Shipan, 2006), and the collective nature of legislative politics and the separation of powers mitigates the influence of any single policy-maker, especially ideologically extreme ones (*e.g.*, Buchanan and Tullock (1962), McCubbins, Noll and Weingast (1987), Krehbiel (1999)). While we do not focus on these complicating factors, we briefly explore theoretical extensions of our model and demonstrate they can actually enhance the strategic dilemma we identify and can under some circumstances exacerbate the politics of scientific expertise.

3 A Model of Political Support for Policy Expertise

We present the simplest model that we have found capable of capturing the incentives of interest.[7] Policy is two-dimensional and all individuals are assumed to have preferences over the two dimensions such that the two dimensions represent complementary goods.[8] This assumption clarifies that we are focused on situations in which there is preference heterogeneity, but not "conflict," *per se*, in the sense that all individuals would prefer higher levels on both issues to lower levels on both issues.[9]

Our theory focuses on two players: a *policy-maker*, P, and an *expert*, E. The policy-maker observes the policy preferences of the expert and then chooses a policy, $z = (x, y)$, where $x \geq 0$, $y \geq 0$, and z satisfies a budget constraint described in more detail below. After observing z, the expert then decides whether to remain in office: he or she will leave office if the policy-maker's choice of policy is too far from the expert's most-preferred policy. Finally, a disaster occurs with some exogenous probability. If the disaster occurs and the expert remained in office, his or her presence mitigates the impact of the disaster.

Policy Preferences. The two dimensions of policy are denoted by x and y, and each player $i \in \{P, E\}$ is characterized by a *type*, $\alpha_i \in [0, 1]$, and his or her payoff from policy $z = (x, y)$ is determined by a Cobb-Douglas payoff function as follows:

$$u_i(z \mid \alpha_i) = x^{\alpha_i} y^{1-\alpha_i}. \tag{1}$$

[7]Specifically, we do not believe that the approach we take in this model (namely, that disasters constrain the set of feasible policies and that experts mitigate this impact) can be incorporated in a unidimensional model with spatial preferences. In such a model, expertise will be more valuable to policy-makers with extreme preferences. We discuss this issue and provide an example of such a model in Appendix G.

[8]In equilibrium, policy-making will be a unidimensional phenomenon, but that is the result of the policy payoff function we adopt for simplicity.

[9]Of course, this is not to say that we believe that all political decisions are characterized by such preferences. Instead, we are demonstrating that fundamentally opposed preferences are not required for the dynamics that our theory uncovers.

Disasters. Our theory is predicated on uncertainty about the future policy environment. Specifically, we assume that, after the policy-maker chooses z and the expert chooses e, a disaster occurs ($d = 1$) with (known) probability $p \in (0,1)$ and, otherwise, no disaster occurs ($d = 0$).

The Budget Constraint. The policy-maker faces a budget constraint that is a function of both whether a disaster occurred (d) and whether the expert remained in office (e). This constraint is denoted by $B(d,e) > 0$ and defined as follows:[10]

$$B(d,e) = 1 - \delta(1 - \kappa e)d,$$

where $\delta \in (0,1)$ represents the budgetary impact of an unmitigated disaster and $\kappa \in (0,1)$ denote the capacity of the expert to mitigate the impact of the disaster when it occurs. In the baseline model, we assume that both δ and κ are exogenous and common knowledge. With this in hand, the set of *feasible policies* depends on $B(d,e)$ as follows:[11]

$$Z(d,e) \equiv \{(x,y) \in \mathbf{R}_+^2 : x + y \leq B(d,e)\}.$$

Before continuing, it is worth noting that the budget constraint need not represent a financial constraint—rather, it simply reflects the fact that in equilibrium increasing policy on either of the dimensions will require reducing policy on the other dimension. Similarly, the effect of a disaster is to tighten this constraint, and this effect is mitigated somewhat with the assistance of experts ($e = 1$).[12] The pandemic unfolding around the world during the writing

[10]The normalization of the disaster-free budget constraint to equal 1 is without loss of generality in our setting. In a more dynamic model, it would be interesting to relax this and allow the total budget in the future be a function of (say) p. This would be particularly relevant for applications to issues such as financial regulation and economic crises.

[11]Our budget constraint presumes that the marginal costs of x and y are equal. It would be simple to extend the model to incorporate different costs for the two dimensions and potentially allow disasters and/or government expertise affect these costs. For reasons of space, we leave such an extension to future work.

[12]For reasons of space, we do not model the actual mechanics of policy-making during a disaster. Rather, we are simply assuming that the policy maker faces a more constrained optimization problem in a disaster, and this constraint is more extreme when the policy-maker does not have the benefit of expertise. This could

of this article provides, sadly, many examples of a disaster constraining the choices available to policy-makers.

Induced Preferences. Each player i's *ideal point* conditional on (d, e) is defined as follows:

$$z_i^*(B(d,e), \alpha_i) = (\alpha_i B(d,e), (1-\alpha_i)B(d,e)).$$

As long as the context is clear, we will write $z_i^*(d,e) \equiv z^*(B(d,e), \alpha_i)$. An example of the policy setting and preferences is displayed in Figure 1.

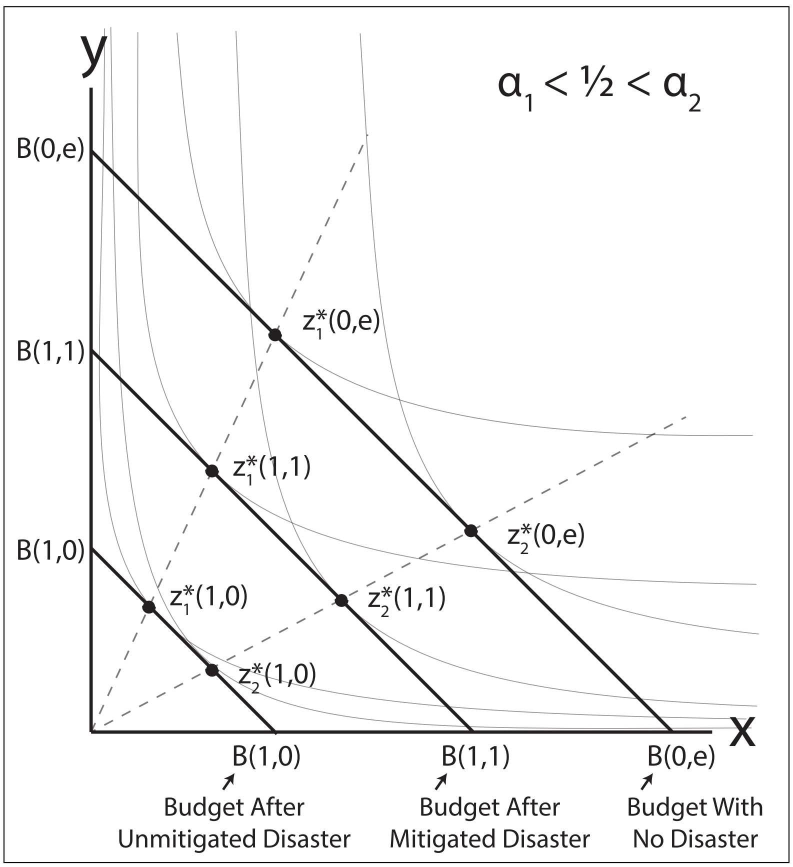

Figure 1: Disasters, Expertise, Feasible Policies, and Induced Ideal Points

Comparing Policy Preferences. As illustrated in Figure 1, Player i's induced ideal point is closer to the y-axis when $\alpha_i < 1/2$, and closer to the x-axis when $\alpha_i > 1/2$. Accordingly,

be because the disaster actually involves the loss of resources or because it requires that resources be used to respond to the disaster (or, probably, both). Either interpretation is consistent with our approach.

we will refer to an individual whose type, α_i, is less than $^1/_2$ as *y-leaning* and an individual whose type is greater than $^1/_2$ as *x-leaning*. We will refer to an individual i as having "more extreme" policy preferences than an individual j if $|\alpha_i - {}^1/_2| > |\alpha_j - {}^1/_2|$.

3.1 Experts and Expertise

Our focus in this article is on the tension between policy change and retaining government expertise. We operationalize this tension by including a unitary "expert" as a policy-motivated actor within the model. We denote this expert by E, and his or her type is denoted by $\alpha_E \in (0, 1)$. We assume that the expert's policy payoff is also defined by Equation 1, and denote the (exogenous and common knowledge) ideal point of the expert by $\alpha_E \in (0, 1)$.

The Costs of Expertise. We assume that retaining expertise in government may require policy concessions (*e.g.*, in the form of granting discretion over policy to "the experts," as in Gailmard and Patty (2007*b*)). In line with this, we assume that setting policy such that the expert's expected policy payoff is strictly less than some exogenous (and known) reservation level, $\rho \geq 0$, will induce the expert to choose $e = 0$ and leave government service.[13] If the policy yields the expert an expected payoff of at least ρ, then he or she chooses $e = 1$ and remains in government service. Thus, the policy-maker may be constrained in adjusting policy if it wants to retain the expert.[14]

[13]For simplicity, we presume that the expert is motivated solely by the policy he or she is charged with implementing. We consider a richer model of motivations in Appendix E.

[14]Note that our assumptions treat ρ and p as independent of one another. From an empirical standpoint, this is arguably unreasonable, assuming that the expert cares about policy outcomes to some degree even if he or she has left government service. Accounting for this in the most reasonable way would require that ρ be a decreasing function of p. Given that our theory already indicates that the policy-maker will be less willing to move policy so far as to induce the expert to leave the government, and making ρ a decreasing function of p would grant the policy-maker greater discretion in terms of how far policy could be moved away from the expert's most preferred allocation without inducing the expert to leave government service, generalizing the model in this way would simply strengthen our key findings, at the expense of greater technical baggage.

Policy-making During a Disaster. If a disaster occurs, the prevailing policy, (x, y), is changed to $(1-\delta(1-\kappa e))(x, y)$.[15] This implies that the "balance" of policy between x and y remains unchanged during a disaster. This is obviously a strong assumption. We believe it is warranted in many circumstances, such as when the balance of (say) service provision has to be set up prior to the disaster. However, even if we relaxed the restriction, as long as the principal's preferences (*i.e.*, α_i for the player i who chooses the policy) remain unchanged, the policy-maker would choose this policy conditional on a disaster occurring, so this assumption does not represent a meaningful constraint on the policy-maker's discretion.[16]

Expected Policy Payoffs. For any agent i with type α_i, the expected payoff from a policy $z = (x, y)$, given expertise $e \in \{0, 1\}$, is

$$EU(z \mid \alpha_i, p, \delta, \kappa, e) = (1-p)u_i(z \mid \alpha_i) + p \cdot u_i((1-\delta(1-\kappa e))z \mid \alpha_i).$$

When the context is clear, we will write this simply as $EU(z \mid \alpha_i, e) \equiv EU(z \mid \alpha_i, p, \delta, \kappa, e)$. Note that $EU(z_i^*(0) \mid \alpha_i, 0)$ is the minimum policy payoff that the policy-maker, potentially constrained only by the expert's decision about whether to remain in government, will receive in equilibrium. It represents the minimum policy payoff that the policy-maker will accept to retain the expert.

Finally, for simplicity, the parameters of the game are referred to as a *situation*, and denoted by $\xi \equiv (\alpha_P, \alpha_E, p, \delta, \kappa)$.

[15]Note that, because each player's ideal point is linear in $B(d, e)$, it is sufficient for us to consider the ratio of x to y that the policy-maker would choose. Regardless of the occurrence of a disaster, this ratio is constant.

[16]In other words, to the degree that we think the policy-maker would want to re-balance policy between x and y after a disaster, we are objecting the policy payoff function defined in Equation 1.

4 Equilibrium Analysis

Recall the sequence of play: first, the policy-maker observes the expert's ideal point, α_E, and then chooses $z = (x, y)$. The expert then observes z and either leaves ($e = 0$) or stays ($e = 1$). After this, a disaster occurs ($d = 1$) with probability p or does not ($d = 0$) with probability $1 - p$. The players then receives their payoffs, and the game concludes. This simple structure is easy to analyze, essentially boiling down to the policy-maker choosing between at most two different policies: the policy-maker's ideal point or the policy that maximizes the policy-makers' payoff among those policies that will induce the expert to choose $e = 1$ and remain in government service.

Optimal Policy-making. After accounting for sequential rationality on the part of the policy-maker, P then has a simple calculation. If P wants to retain the expert, then P's optimal policy choice is

$$z_P^R(\alpha_E, \rho) = \underset{z \in \{Z(0,1): EU(z \mid \alpha_E, 1) \geq \rho\}}{\operatorname{argmax}} EU(z \mid \alpha_P, 1).$$

If the policy-maker is unconcerned about retaining the expert, the optimal policy is simply $z_1^*(0)$. Thus, the optimal policy is either z_P^R or $z_1^*(0)$. If ρ is sufficiently small or α_E is sufficiently close to α_P, these two policies will be identical, and the policy-maker is unconstrained by E. That said, the interesting cases will be exactly those in which $z_P^R \neq z_1^*(0)$, in which case the policy-maker is constrained by E in the sense that the policy-maker must choose a policy other than his or her ideal point in order to retain the expert. When the policy-maker is constrained by E in this way, the policy-maker's optimal choice is

$$z^*(p, \delta, \kappa, \alpha_E, \rho) = \begin{cases} z_1^*(0) & \text{if } EU(z_1^*(0) \mid \alpha_P, p, \delta) \geq EU(z_P^R(\alpha_E, \rho) \mid \alpha_P, 1), \\ z_P^R(\alpha_E, \rho) & \text{otherwise.} \end{cases}$$

4.1 What Kinds of Policy-makers Defer to Expertise?

We first focus on the theory's predictions regarding what types of policy-makers will defer to expertise and, correspondingly, what types of experts they will defer to. The key results are that policy-makers with extreme policy preferences are less likely to defer to experts, regardless of their leaning, and policy-makers are more likely to defer to experts who share their policy leaning.

Extremist Policy-makers Undermine Expertise

The first conclusion from the theory is that more extreme policy-makers are more likely to set policy in such a way as to undermine expertise. [17]

Prediction 1 *Policy-makers with more extreme policy preferences will*

1. *be more likely to undermine expertise in pursuit of policy change,*
2. *have a lower expected level of retained government expertise, and*
3. *experience larger expected losses from disasters.*

Note that Prediction 1 is independent of the lean of the policy-maker's policy preferences. By construction, policy-makers with extreme x-leaning preferences and those with similarly extreme y-leaning preferences view the world in opposite but symmetric ways. Relatedly our theory predicts that *any* expert, regardless of his or her preferences, will be replaced by a policy-maker with sufficiently extreme policy preferences. In other words, no expert is "safe" from being undermined by the policy-maker.

In addition, this insecurity is not completely imbalanced in ideological terms: any expert might be replaced by either an x-leaning or y-leaning policy-maker. We will see below that

[17] This is stated and proven formally in Appendix B (Proposition 1). In order to keep the presentation of the results in the body of the article as accessible as possible, we refer to one policy-maker type, α_P as being "more likely" to induce an expert to leave than another policy-maker type, $\alpha_{P'}$, if the set of expert types that would lead P to set policy so as to induce the expert to leave has a larger Lebesgue measure than the set of expert types that would lead P' to set policy so as to induce the expert to leave.

an x-leaning expert is less likely to be replaced by a randomly drawn x-leaning policy-maker than by a randomly drawn y-leaning policy-maker (and vice-versa), but the threat of removal is positive for any expert, regardless of the alignment of the leanings of the expert and the policy-maker.

In this way, the model provides a nuanced understanding of the role of ideology/policy priorities in deference to expertise: observing an x-leaning policy-maker removing an x-leaning expert does not imply that the direction of ideology/policy priorities are irrelevant to deference. Rather, it simply demonstrates that it is not singularly dispositive.

Policy-makers Tend to Undermine Experts With Opposite Policy Leanings

As the policy-maker becomes more extreme in a given direction, he or she finds it in his or her interest to sacrifice expertise in order to achieve policy change for a larger proportion of potential status quo policies. In addition, as the policy-maker becomes more extreme, he or she becomes more accepting of some relatively extreme experts who share the policy-maker's leaning with respect to the balance of x and y, and less accepting of those whose leanings are either less extreme than or opposed to that of the policy-maker. These comparative statics are illustrated in Figure 2 and lead to the following empirical prediction.[18]

Prediction 2 *The policy leanings of the experts retained by the policy-maker will be positively correlated with those of the policy-maker.*

An implication of Prediction 2 is that the governments' experts will be more biased when the policy-maker has more extreme preferences.[19] This prediction can also be interpreted in terms of the appearance of political responsiveness: experts who remain in equilibrium will tend to share the views of the policy-maker. This raises the question of how elections that lead to changes in the policy-maker's ideal point will affect governmental expertise and preparation for disasters.

[18] This prediction is based on Propositions 3 and 4, each of which is stated and proved in Appendix B.

[19] Formally, the ratio of x-leaning types of experts who would remain in government to y-leaning expert who would remain will be farther from $1/2$ when the policy-maker's ideal point is farther from $1/2$.

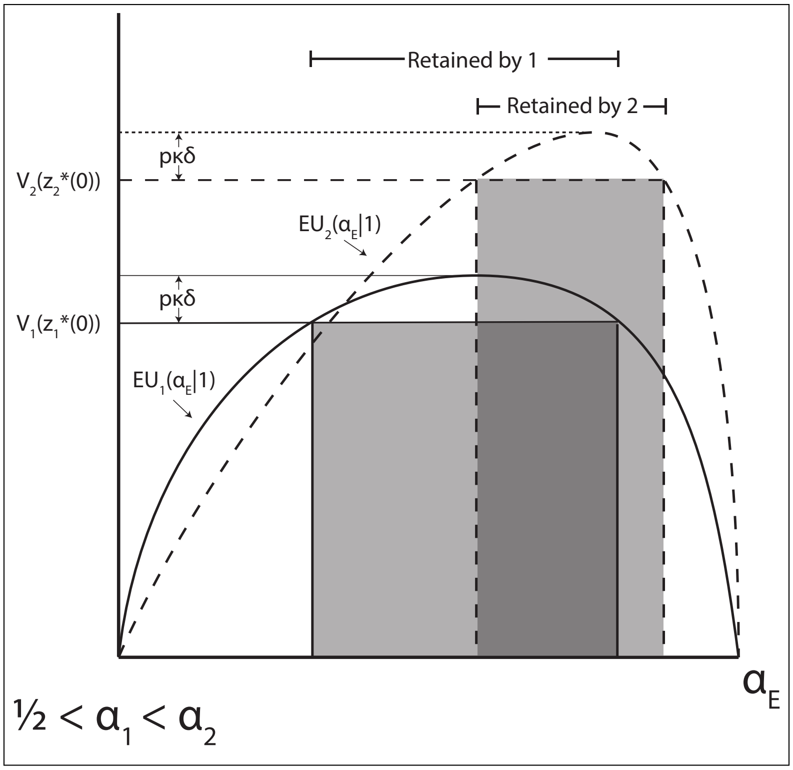

Figure 2: Deference to Experts As a Function of Policy-maker's Ideal Point, α_P

4.2 Electoral Impacts on Deference to Expertise

If we think of the policy-maker's ideal point, α_P, as being drawn from a distribution, our theory offers some predictions about the effect of polarization and ideological shifts on retention of government expertise and the government's response to disasters. Our notions of polarization and electoral shifts are illustrated in Figure 3.[20] With these conceptions of the effect of electoral polarization/shifts in hand, Prediction 1 leads directly to an ancillary prediction regarding the effect of polarization on deference to expertise: as polarization increases, every expert is "eventually" at increased risk of being undermined.[21]

Prediction 3 *The probability that any given expert will be replaced is increasing in the polarization of policy-makers' preferences for sufficiently high levels of polarization.*

[20]The notions are formally defined in Appendix A, and we discuss how this conception of electoral effects can be linked with the much bigger question of electoral accountability in Appendix F.

[21]This prediction follows from our definition of increased polarization in Appendix A and Proposition 2 in Appendix B.

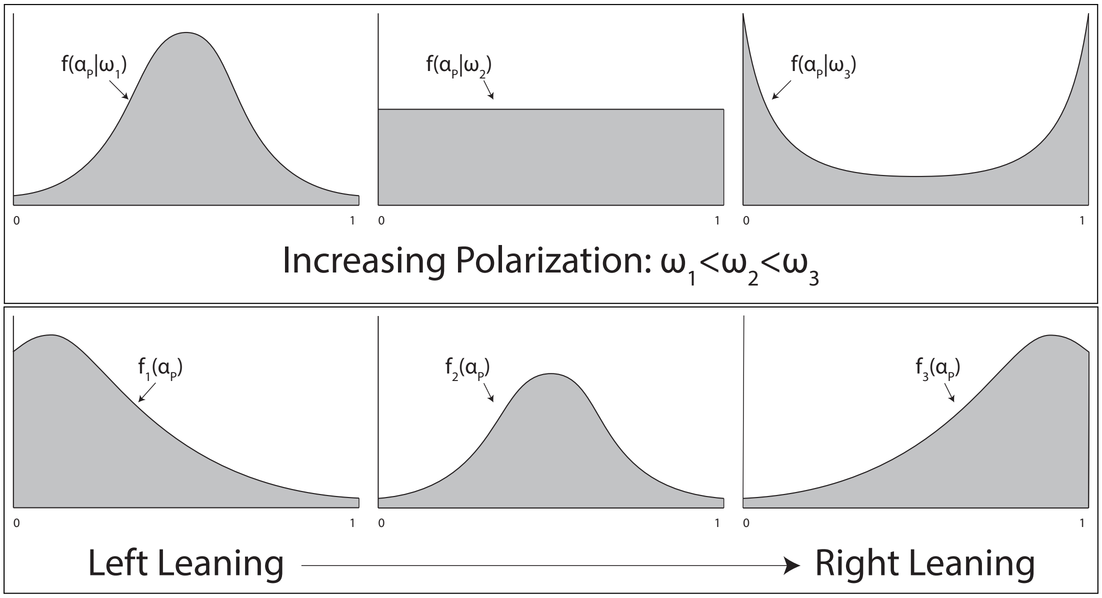

Figure 3: Representing Polarization and Shifts

5 Governmental Failures and Ideological Alignment

Our theory offers a clear operationalization of a *governmental failure*. Specifically, we define a governmental failure as happening whenever a disaster occurs *and* the expert has left the government (*i.e.*, when $d = 1$ and $e = 0$). This is a particularly clean definition for our purposes, because our theory has assumed that the occurrence of a disaster is independent of the expert's presence.[22]

With this in hand, the final set of theoretical conclusions we discuss in the body of the article complete the circle by summarizing the implications for the correlation between the occurrence of governmental failure and the extremism of the policy-maker's preferences. An illustration of the distribution of the policy-maker's preferences, conditional on the occurrence of government failure is displayed in Figure 4.

In the example pictured in Figure 4, the preferences of the policy-maker and expert, α_P and α_E, are assumed to be independently and uniformly distributed. Accordingly, the

[22]If we assumed, as seems reasonable, that the expert leaving government increases the risk of a disaster occurring, then the conclusions below would be even stronger, because that possibility would increase the range of experts who would be retained by any policy-maker, thereby exacerbating the proportion of policy-makers with extreme preferences among those who undermine expertise in equilibrium.

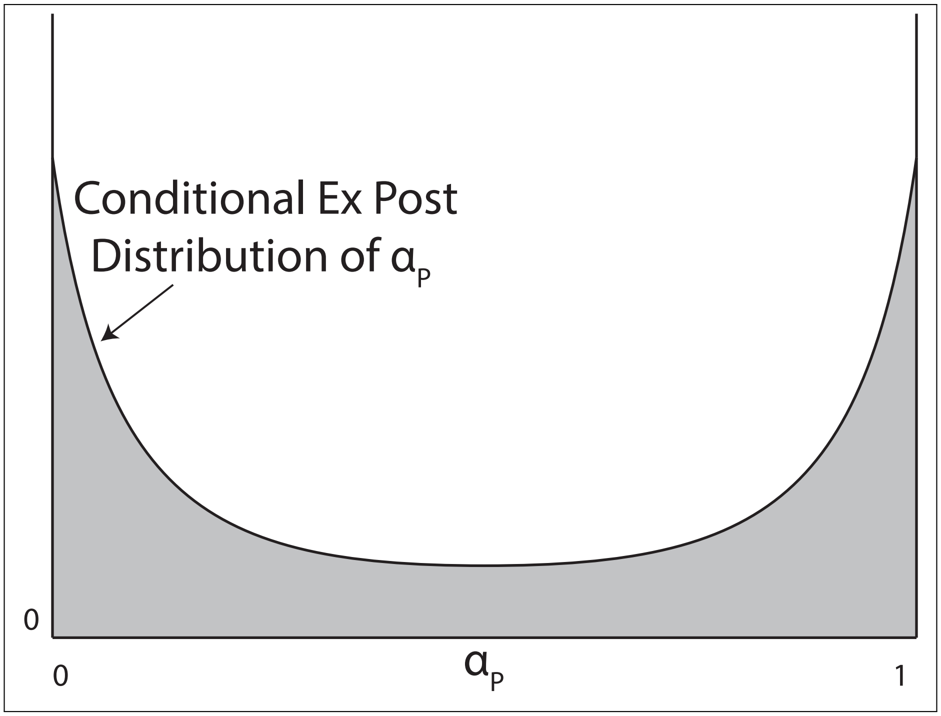

Figure 4: Distribution of the Policy-maker's Preferences, Conditional on Government Failure

conclusion to be drawn from the figure is that, because the occurrence of a disaster is correlated with the realized preference of the policy-maker (*i.e.*, the election result), policy-makers with extreme preferences will be more likely to be responsible for a government failure.

If, on the other hand, we assume that the distribution of the expert's preferences are biased toward (say) x-leaning preferences, then the analogue of Figure 4 is illustrated in Figure 5. The figure indicates that, conditional on government failure, the policy-maker will tend to not only have extreme preferences, but also tend to have preferences that are of the opposite leaning relative to any ex ante bias in the leaning of the government experts.

In terms of application, an informal implication of Figure 5 is that governmental failures should be more common after a dramatic change in the policy goals of the policy-maker. After all, the experts retained by a fairly y-leaning policy-maker would be exactly those that a sufficiently x-leaning policy-maker would not defer to, and vice-versa. This is in line with the fact that governmental failures are correlated with increased turnover within

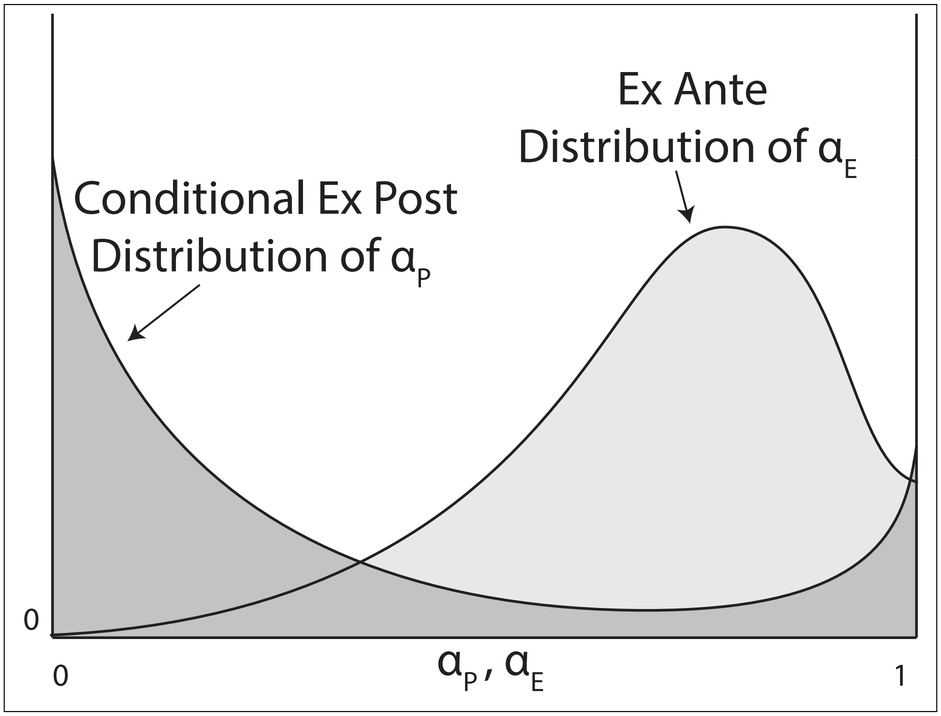

Figure 5: Policy-maker's Preferences, Conditional on Failure, With Biased Distribution of Experts

the relevant government agencies, and increased turnover is correlated with reductions in discretion and/or resources allocated to an agency (Brehm and Gates (1997), Bertelli (2007)).

The Politics of Expertise in the US. The analysis in this section indicates why one might expect that "small government" candidates who implement platforms intended to limit the discretion and policy reach of government agencies and their employees are more associated with government failure in this framework—it is not necessarily that such candidates do not respect expertise but, rather, that their policy preferences value policies that experts who have chosen careers in the government do not favor. It is crucial to remember that, in our framework, the policy-maker *always* strictly prefers retaining the expert to the expert leaving government service—*holding the implemented policy constant.* In other words, the theoretical framework is designed to favor retention of expertise—the loss of expertise (and accordingly, the possibility of governmental failure) is driven entirely by policy preferences rather than any innate distaste for expertise, per se.

6 Extensions

Collegial Policy-making. The theory presented above presumes, for simplicity, that policy is set in a unilateral fashion. In Appendix C, we extend the model to allow for collective choice in both majoritarian and super-majoritarian settings. To the degree that the ideal points of the multiple policy-makers in such a setting are at least somewhat independent of one another, equilibrium policy-making will tend to choose more centrist policies and, accordingly, be more likely to retain an expert with a randomly drawn ideal point. Similarly, this "expertise protection" effect becomes more pronounced as the threshold of support required for policy change becomes more demanding.

Uncertainty About the Expert's Effectiveness. The theory presented above assumes that the expert's effectiveness in mitigating the effects of a disaster (*i.e.*, κ) is common knowledge between the policy-maker and expert. In Appendix D, we relax this assumption somewhat by allowing the policy-maker to observe a noisy signal about κ prior to setting policy. One interest in that extension is the conditions under which the policy-maker's decision regarding whether to retain expertise will be responsive to the policy-maker's signal about κ.

We show (Proposition 7) that deference to the expert will be responsive to the policy-maker's information only if the expert's policy preferences are different enough from those of the policy-maker and the signal conveys sufficiently large information (in terms of moving the policy-maker's posterior beliefs about κ).

These are both intuitive conclusions, of course, but a subtle aspect of the extension is that the level of information conveyed by the signal required for the policy-maker to respond to it in his or her policy choice is an increasing function of the dissimilarity between the expert's and policy-maker's preferences.

Finally, the policy-maker should be responsive to his or her signal about κ when the expected impact of the disaster in question is moderate (*i.e.*, $p \cdot \delta$ is neither too close to zero nor too close to one). This conclusion is informative regarding the types of contingencies that the policy-maker should expend resources investigating (for example, through planning within agencies and other efforts such as information collecting by institutional organs such as the Government Accountability Office in the United States federal government).

7 Discussion and Conclusion

Among the core functions of government is providing organized, large-scale responses to major disasters, especially when typical market failures make private responses inadequate. Natural disasters, health pandemics, and other catastrophes threaten to security, prosperity, and safety of people, and they look to their government to provide effective and decisive responses when threatened with these events. However, by their very nature, these kinds of events are unpredictable and require policy and action that is atypical. As a consequence, responses to disasters are more effective when they benefit from technical or scientific expertise. Details and nuance can make the difference when the stakes are high and the relationship between policies and outcomes is hard to know. Indeed, that knowledge is precisely what expertise provides to policy-makers.

However, in order to take advantage of scientific experts, politicians must retain them in government service, so that they are available to provide technical knowledge when disaster strikes. Donald Trump's efforts to "drain the swamp" after his election had the effect, for example, of inducing scientists and experts to leave government service, deteriorating the base of scientific capacity on which the government could draw in times of crisis. As a result, in 2020, there is still no scientific consensus on the safety of drinking water in Flint, Michigan, and during the covid-19 pandemic, government responses have been inconsistent, slow, and inadequate. In one now-infamous quote from 2018, President Trump said, in defense of cuts to funding for the US pandemic response team,

> We can get money, we can increase staff—we know all the people. This is a

> question I asked the doctors before. Some of the people we cut, they haven't used for many, many years, and if we have ever need them we can get them very, very quickly. And rather than spending the money—I'm a business person. I don't like having thousands of people around when you don't need them. When we need them, we can get them back very quickly.[23]

Arguably, President Trump might havehonestly approaching the federal bureaucracy as a "business person" with a high tolerance for risk. However, that he believed that "[w]e can get money, we can increase staff... we can get them very, very quickly." However, what scholars of government, and indeed leaders with experience in the government, know is that it is not possible to rapidly develop bureaucratic knowledge, capacity, and expertise. In particular, when disaster, such as pandemics, strike, government is called upon to rapidly respond, and to do so requires the establishment of a readily-available infrastructure of technical capacity.

What is more, it is also possible—and not necessarily contradictory—that President Trump's approach to cutting public health funding was not driven by an honest, if mistaken, business-person approach to bureaucratic expenditures. It is possible that he was also motivated by ideological conflict with public health agencies. Having trafficked in conspiracy theories, President Trump embraced opponents of vaccines, spreading debunked claims that vaccinations have been linked to autism.[24] (Though, notably, Trump shifted his position modestly in 2019 on he tails of a measles outbreak that has ostensibly been facilitated by a decline in childhood vaccinations in the US.)

We have argued that, in general, these two motivations induce politicians to undermine their own access to scientific and technical expertise. To develop and retain scientific capacity, politicians must adopt agendas and policies that are attractive enough to scientific experts, and they must believe that the constraining effect of that tactic is justified relative to their view of the potential benefit of expertise. By modeling these incentives, we have shown that

[23] https://www.independent.co.uk/news/world/americas/us-politics/coronavirus-video-trump-pandemic-team-cut-2018-a9405191.html

[24] https://www.statnews.com/2019/04/26/trump-vaccinations-measles/

failed government responses to disasters, such as pandemics, are more likely to arise in the wake of leadership by ideological extremists. It is those leaders who will have the strongest incentive to selectively undermine scientific capacity, leading to a deterioration of technical experts among the bureaucratic workforce. *Polarization about pandemics arises because they become disastrous precisely when our politics are already polarized.*

Stepping back, this theoretical expectation helps make sense of the sometimes sense relationship between policy-makers and their advisors. Sometimes, a significant concern for politicians is facing a decision between doing what is right for her constituents and doing what her constituents want. Under the best of circumstances, a leader knows what is best, and it aligns with his political interest. Under less fortunate situations, those forces push in opposite directions, and democratic goals of accountable governance and responsible governance can come into conflict. In these settings, there is a very real risk of institutional pathologies undermining effective leadership. A sensible expectation might be, then, that a politician torn between doing what is right and doing what is popular might take advantage of technical expertise to shield herself from reprisal. Unfortunately, this is not always the case, and we sometimes observe politicians attacking the very experts who might help them make better decisions or provide political cover for unpopular policies that are nevertheless in the public interest.

In light of our analysis, one reason why we might not see politicians relying on scientific expertise in order to take unpopular decisions is that they might not have access to that expertise. In this paper, we have explored one mechanism by which this phenomenon arises—politicians sometimes make choices that reduce their access to scientific expertise. In particular, we have modeled the effects of ideological extremism that reduces the attractiveness of government work to technical experts. Underlying our analysis is an assumption that scientific experts have their own preferences about policy and as a consequence there are policies a political leader might adopt that lead the expert to quit her job. Simply put, scientific experts care about the policies for which their expertise is used to implement. As we have shown, a politician might be keenly aware of this dynamic, and as a consequence ideological extremists might have an incentive to actively "chase off" experts, relative to more moderate politicians. They might do so specifically because the policy concessions necessary to retain the expert are not worth the value the expert brings in terms of mitigating the risk of a disaster.

To understand why we so often see this counter-intuitive behavior, it is instructive to consider how expertise operates in politics. The study of policy-making is replete with examples of the value of information. For example, scholars have extensively examined ways in which information can improve the effectiveness of policy-making. Related work has focused on strategic tensions faced when policy-makers are faced with a trade-off between taking advantage of substantive expertise and ensuring ideological alignment with those developing that expertise. In this spirit, a wide class of theoretical models examine the conditions that facilitate efficient transmission of information through a policy-making process. One of the canonical lessons from this line of inquiry is that policy expertise often comes at a cost of giving up some control or binding one's degree of ultimate discretion. In an important sense, this lesson captures the foundational difficulty in politics that is captured by the ally principle typical of principal-agent settings.

In the setting of national disasters or health pandemics, the stakes are typically raised, compared to normal policy-making needs. At the same time, by their very nature, the risk of a global health pandemic or major disaster arising is typically low, and these events are typically not foreseen. However, it is precisely in these settings—where unusual, unpredictable, disastrous events arise—that scientific or technical expertise is *most* valuable to a policy-maker. But, our analysis shows that it was ideological bias *prior* to the pandemic that led us to a situation in which we are unprepared to respond adequately. Therefore, when we see conflict between policy-makers, it is likely to be either in the context of scientific expert opinion constraining a policy-maker in the absence of a national emergency or, alternatively, a national emergency emerging under conditions where scientific expertise is lacking because of prior hostilities.

National emergencies, including natural disasters and health pandemics, are events that stress a government's capacity. They are also moments that ostensibly align interests among the public, lowering traditional distributive battles in favor of protecting health and safety. However, it is not uncommon to see sharp partisan polarization around these events. The intricacies of the politics of bureaucratic expertise and capacity are not uniquely ideological in the context of national emergencies, but their implications are perhaps sharpest. It is for this reason that, though disasters are typically out of mind during times of normal politics, it is perhaps in that context that policy-makers need to most thoughtfully consider how and when to establish and insulate scientific capacity in governance.

References

Austen-Smith, David and Jeffrey S. Banks. 1996. "Information Aggregation, Rationality and the Condorcet Jury Theorem." *American Political Science Review* 90:3445.

Bertelli, Anthony M. 2007. "Determinants of bureaucratic turnover intention: Evidence from the Department of the Treasury." *Journal of Public Administration Research and Theory* 17(2):235–258.

Bertelli, Anthony M and Christian R Grose. 2009. "Secretaries of pork? A new theory of distributive public policy." *The Journal of Politics* 71(3):926–945.

Bertelli, Anthony M and David E Lewis. 2012. "Policy influence, agency-specific expertise, and exit in the federal service." *Journal of Public Administration Research and Theory* 23(2):223–245.

Bolton, Alexander, John M De Figueiredo and David E Lewis. 2016. Elections, ideology, and turnover in the US federal government. Technical report National Bureau of Economic Research.

Brehm, John and Scott Gates. 1997. *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public.* Ann Arbor, MI: University of Michigan Press.

Buchanan, James M. and Gordon Tullock. 1962. *The Calculus of Consent: Logical Foundations of Constitutional Democracy.* Ann Arbor, MI: University of Michigan Press.

Cameron, Charles M and John M de Figueiredo. 2018. Quitting in protest: A Theory of Presidential policy making and civil service response. Technical report Working paper.

Carpenter, Daniel and George A Krause. 2014. "Transactional Authority and Bureaucratic Politics." *Journal of Public Administration Research and Theory* 25(1):5–25.

Clinton, Joshua D and David E Lewis. 2008. "Expert opinion, agency characteristics, and agency preferences." *Political Analysis* pp. 3–20.

Dewatripont, Mathias, Ian Jewitt and Jean Tirole. 1999. "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies." *The Review of Economic Studies* 66(1):199–217.

Doherty, Kathleen M, David E Lewis and Scott Limbocker. 2019. "Executive control and turnover in the senior executive service." *Journal of Public Administration Research and Theory* 29(2):159–174.

Gailmard, Sean and John W. Patty. 2007*a*. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51(4):873–889.

Gailmard, Sean and John W. Patty. 2007*b*. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51(4):873–889.

Gailmard, Sean and John W. Patty. 2012. *Learning While Governing: Information, Accountability, and Executive Branch Institutions.* Chicago, IL: University of Chicago Press.

Gordon, Sanford C. 2011. "Politicizing agency spending authority: Lessons from a Bush-era scandal." *American Political Science Review* pp. 717–734.

Hofstadter, Richard. 1963. *Anti-intellectualism in American Life.(3. Print.).* Vol. 713 Vintage Books.

Huber, John D and Charles R Shipan. 2006. Politics, delegation, and bureaucracy. In *The Oxford handbook of political science.*

Huber, John D., Charles R. Shipan and Madelaine Pfahler. 2001. "Legislatures and Statutory Control of Bureaucracy." *American Journal of Political Science* 45(2):330–345.

Jacoby, Susan. 2018. *The age of American unreason in a culture of lies.* Vintage.

Krause, George A and Anne Joseph O'Connell. 2019. "Loyalty–Competence Trade-offs for Top US Federal Bureaucratic Leaders in the Administrative Presidency Era." *Presidential Studies Quarterly* 49(3):527–550.

Krehbiel, Keith. 1998. *Pivotal Politics: A Theory of U.S. Lawmaking.* Chicago, IL: University of Chicago Press.

Krehbiel, Keith. 1999. "Paradoxes of Parties in Congress." *Legislative Studies Quarterly* 24:31–64.

Lewis, David E. 2009. "Revisiting the Administrative Presidency: Policy, Patronage, and Agency Competence." *Presidential Studies Quarterly* 39(1):60–73.

McCubbins, Mathew, Roger Noll and Barry Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law, Economics, and Organization* 3(2):243–277.

Moe, Terry M. 1985. The Politicized Presidency. In *The New Direction in American Politics*, ed. John Chubb and Paul Peterson. Washington, DC: Brookings Institution.

Monroe, Nathan, John W Patty and Elizabeth Maggie Penn. 2018. "What Is Pivotal Politics (and What Else Can It Be)?" *The Journal of Politics* 80(3):1088–1099.

Nelkin, Dorothy. 1975. "The political impact of technical expertise." *Social studies of science* 5(1):35–54.

Patty, John W., James M. Snyder and Michael M. Ting. 2009. "Two's Company, Three's An Equilibrium: Strategic Voting and Multicandidate Elections." *Quarterly Journal of Political Science* 4:251–278.

Perry, James L. and Lois Recascino Wise. 1990. "The Motivational Bases of Public Service." *Public Administration Review* 50(3):367–373.

Richardson, Mark D, Joshua D Clinton and David E Lewis. 2018. "Elite perceptions of agency ideology and workforce skill." *The Journal of Politics* 80(1):303–308.

Rudalevige, Andrew. 2009. "The administrative presidency and bureaucratic control: implementing a research agenda." *Presidential Studies Quarterly* 39(1):10–24.

Rudalevige, Andrew and David E Lewis. 2005. Parsing the politicized presidency: Centralization and politicization as presidential strategies for bureaucratic control. In *Annual Meeting of the American Political Science Association, Washington, DC.*

Snyder, Susan K and Barry R Weingast. 2000. "The American system of shared powers: the President, Congress, and the NLRB." *Journal of Law, Economics, and Organization* 16(2):269–305.

A Modeling Polarization and Bias

When relevant, we assume that policy-maker's ideal point is distributed according to a distribution governed by a cumulative distribution function, F. We assume that F is continuously differentiable and assigns positive measure to all open sets in $[0,1]$[25] and then capture ideological polarization and ideological bias in the electoral process by parameterizing this cumulative distribution function as follows.

Polarization. We will refer to policy-making as being more *polarized* when the tails of the distribution get "fatter." In order to represent this precisely, let $\mathcal{F} : \mathbf{R}_{++} \to \Delta([0,1])$ be a family of distributions such that, for any $\omega > 0$, the cumulative distribution function is denoted by $F(\cdot \mid \omega)$, and the family $\mathcal{F}$ satisfies the following:

- $F({}^1\!/_2 \mid \omega) = {}^1\!/_2$,
- $F(\alpha_P \mid \omega) = 1 - F(1 - \alpha_P \mid \omega)$ for all α_P,
- $\frac{\partial F(\alpha_P \mid \omega)}{\partial \omega} > 0$ for $\alpha_P < {}^1\!/_2$, and
- $\frac{\partial F(\alpha_P \mid \omega)}{\partial \omega} < 0$ for $\alpha_P > {}^1\!/_2$.

In this setting, increasing polarization is equivalent to increasing ω. An example of such a family of distributions is the $\text{Beta}(\omega^{-1}, \omega^{-1})$ family.

Ideological Bias. We will refer to policy-making as becoming *more right-leaning* when the new distribution of the policy-maker's type first-order stochastically dominates the old distribution, and conversely, as *more left-leaning* when the old distribution first-order stochastically dominates the new.

[25] This regularity condition is equivalent to assuming $F(x) \in (0,1)$ for all $x \in (0,1)$.

B Formal Results and Proofs

Most of our formal results revolve around two related sets, describing in turn the set of policy-maker types that will retain a given expert type and, conversely, the set of expert types that will be retained by a given policy-maker type.

For any given setting $(p, \delta, \kappa, \alpha_E, \rho)$, the set of policy-maker types who will defer to the expert and choose z_P^R (which might be equal to $z_1^*(0)$) is denoted by

$$M(p, \delta, \kappa, \alpha_E, \rho) = \left\{ \alpha_P \in [0,1] : EU(z_P^R(\alpha_E, \rho) \mid \alpha_P, 1) > EU(z_1^*(0) \mid \alpha_P, p, \delta) \right\}.$$

Similarly, for any given $(p, \delta, \kappa, \alpha_P, \rho)$, the set of Expert types that the policy-maker would choose to retain (*i.e.*, choose $z_P^R(\alpha_E, \rho)$) is denoted by

$$R(\alpha_P, p, \delta, \kappa, \rho) = \left\{ \alpha_E \in [0,1] : EU(z_P^R(\alpha_E, \rho) \mid \alpha_P, 1) \geq EU(z_1^*(0) \mid \alpha_P, p, \delta) \right\}.$$

Extremist Policy-makers Undermine Expertise. The first conclusion from the theory is that any expert will be replaced by a policy-maker with sufficiently extreme policy preferences.

Proposition 1 *For any given setting* $(p, \delta, \kappa, \alpha_E, \rho)$*,* $M(p, \delta, \kappa, \alpha_E, \rho)$ *is an interval strictly contained in* $(0, 1)$*.*

Proof: First, $M(p, \delta, \kappa, \alpha_E, \rho) \neq \varnothing$ because $\alpha_E \in M(p, \delta, \kappa, \alpha_E, \rho)$. $EU(z_1^*(0) \mid \alpha_P, p, \delta)$ is a strictly concave function of α_P, establishing that $M(p, \delta, \kappa, \alpha_E, \rho)$ is an interval. Finally,

$$\lim_{\alpha_P \to 0} EU(z_P^R(\alpha_E, \rho) \mid \alpha_P, 1) = \lim_{\alpha_P \to 1} EU(z_P^R(\alpha_E, \rho) \mid \alpha_P, 1) = 0,$$

and

$$EU(z_1^*(0) \mid \alpha_P, p, \delta) > 0$$

for all α_P, establishing that this interval is strictly contained within the $[0, 1]$ interval. ∎

The next proposition states that the size of the set of expert ideal points that a policy-maker will retain in equilibrium is decreasing in the extremity of the policy-maker's ideal point, α_P. While the result is stated in terms of $\alpha_P > 1/2$, the result is symmetric for $\alpha_P < 1/2$.

Proposition 2 *Suppose that $\alpha_{P'} > \alpha_P > 1/2$. Then for any $(p, \delta, \kappa, \rho)$, $R(\alpha_{P'}, p, \delta, \kappa, \rho)$ is narrower than $R(\alpha_P, p, \delta, \kappa, \rho)$:*

$$\int_{R(\alpha_{P'}, p, \delta, \kappa, \rho)} 1 \ \mathrm{d}\alpha_E < \int_{R(\alpha_P, p, \delta, \kappa, \rho)} 1 \ \mathrm{d}\alpha_E.$$

Proof: Fix $(p, \delta, \kappa, \rho)$ and let $\bar{z}_1(\alpha) < \alpha < \bar{z}_2(\alpha)$ denote the two indifference points for the policy-maker with $\alpha_P = \alpha$. These are defined by the following:

$$\begin{aligned} EU(\bar{z}_1(\alpha)) = EU(\bar{z}_2(\alpha)) &= (1-p)\alpha^\alpha(1-\alpha)^{1-\alpha} + p((1-\delta)\alpha)^\alpha((1-\delta)(1-\alpha))^{1-\alpha}, \\ &= (1-p\delta)\alpha^\alpha(1-\alpha)^{1-\alpha}, \end{aligned}$$

Without loss of generality and to simplify the analysis, we focus on the normalized policy payoff function, $\frac{EU(x)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}$. This implies that $\bar{z}_1(\alpha)$ and $\bar{z}_2(\alpha)$ are defined as follows:

$$\begin{aligned} \frac{EU(\bar{z}_1(\alpha))}{\alpha^\alpha(1-\alpha)^{1-\alpha}} = \frac{EU(\bar{z}_2(\alpha))}{\alpha^\alpha(1-\alpha)^{1-\alpha}} &= (1-p)\alpha^\alpha(1-\alpha)^{1-\alpha} + p((1-\delta)\alpha)^\alpha((1-\delta)(1-\alpha))^{1-\alpha}, \\ &= 1-p\delta. \end{aligned}$$

Then, writing $\tau(x) \equiv \tanh^{-1}(x)$ and noting that $\tau(x)$ inherits the sign of x and $|\tau(x)|$ is symmetric about zero, the marginal effect of α on any given policy $(x, 1-x)$ is

$$\begin{aligned} \frac{\partial}{\partial\alpha}\left[\frac{EU(x)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}\right] &= -(1-\alpha)^{\alpha-1}\alpha^{-\alpha}(1-x)^{-\alpha}x^{\alpha-1}\left((x-\alpha)x' + 2(x-1)x\left(\tau(1-2\alpha)-\tau(1-2x)\right)\right), \\ &= 2(1-\alpha)^{\alpha-1}\alpha^{-\alpha}(1-x)^{1-\alpha}x^\alpha\left(\tau(1-2\alpha)-\tau(1-2x)\right) \end{aligned}$$

For any $\alpha \in [0,1]$ and $\delta > 0$, let $x_L(\alpha, \delta) \equiv \alpha - \delta$ and $x_R(\alpha, \delta) \equiv \alpha + \delta$ denote the two equidistant points $\delta > 0$ from α. The result follows from the fact that the marginal effect of changing α on the normalized policy payoff function is larger in magnitude for x_L when $\alpha > 1/2$ and larger for x_R when $\alpha < 1/2$. To see this, note

$$\begin{aligned}
\frac{\left|\frac{\partial}{\partial\alpha}\left[\frac{EU(x_L)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}\right]\right|}{\frac{\partial}{\partial\alpha}\left[\frac{EU(x_R)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}\right]} &= \frac{|2(1-\alpha)^{\alpha-1}\alpha^{-\alpha}x_L^\alpha(1-x_L)^{1-\alpha}\left(\tau(1-2\alpha)-\tau(1-2x_L)\right)|}{2(1-\alpha)^{\alpha-1}\alpha^{-\alpha}x_R^\alpha(1-x_R)^{1-\alpha}\left(\tau(1-2\alpha)-\tau(1-2x_R)\right)} \\
&= \frac{2(1-\alpha)^{\alpha-1}\alpha^{-\alpha}x_L^\alpha(1-x_L)^{1-\alpha}\left(\tau(1-2x_L)-\tau(1-2\alpha)\right)}{2(1-\alpha)^{\alpha-1}\alpha^{-\alpha}x_R^\alpha(1-x_R)^{1-\alpha}\left(\tau(1-2\alpha)-\tau(1-2x_R)\right)}, \\
&= \frac{x_L^\alpha(1-x_L)^{1-\alpha}}{x_R^\alpha(1-x_R)^{1-\alpha}},
\end{aligned}$$

so that

$$\begin{aligned}
\alpha < 1/2 &\Rightarrow \frac{\left|\frac{\partial}{\partial\alpha}\left[\frac{EU(x_L)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}\right]\right|}{\frac{\partial}{\partial\alpha}\left[\frac{EU(x_R)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}\right]} < 1, \text{ and} \\
\alpha < 1/2 &\Rightarrow \frac{\left|\frac{\partial}{\partial\alpha}\left[\frac{EU(x_L)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}\right]\right|}{\frac{\partial}{\partial\alpha}\left[\frac{EU(x_R)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}\right]} > 1,
\end{aligned}$$

which implies that the first derivative of the point closest to $1/2$ at which $\frac{EU(x)}{\alpha^\alpha(1-\alpha)^{1-\alpha}}$ intersects the horizontal line $y = 1 - p\delta$ has a larger magnitude than that of the point farthest from $1/2$.[26] This implies that as was to be shown. This implies that the length of the interval $R(\alpha, p, \delta, \kappa, \rho)$ is increasing in $\alpha \in [0, 1/2)$ and decreasing in $\alpha \in (1/2, 1]$, so that $\alpha_{P'} > \alpha_P > 1/2$ implies

$$\int_{R(\alpha_{P'},p,\delta,\kappa,\rho)} 1 \ \mathrm{d}\alpha_E < \int_{R(\alpha_P,p,\delta,\kappa,\rho)} 1 \ \mathrm{d}\alpha_E,$$

as was to be shown. ∎

Policy-makers Tend to Undermine Experts With Opposed Policy Preferences. The next conclusion from the baseline model is that policy-makers will, in a specific sense, be more likely to remove experts whose policy preferences have the opposite leaning of those of the policy-maker. In other words, the model offers a prediction of "ideologically biased"

[26] And, it is easily verified, the first derivative of each of these points is strictly positive with respect to $\alpha \in (0, 1)$.

losses of expertise.

Proposition 3 *If the policy-maker is y-leaning, then the Lebesgue measure of the set of ideal points of x-leaning experts that the policy-maker would remove is larger than that of the set of ideal points of y-leaning experts that he or she would remove (and vice-versa).*

Proof: Note that, because $p > 0$ and $\delta > 0$, $R(\alpha_P, p, \delta, \kappa, \rho)$ is an interval strictly containing α_P. Denote the infimum and supremum of $R(\alpha_P, p, \delta, \kappa, \rho)$ by r_0 and r_1 respectively. When $\alpha_P = {}^1\!/_2$, $r_0 = 1 - r_1$, so that the Lebesgue measure of x-leaning experts that the policy-maker would remove is equal to that of the y-leaning experts he or she would remove when $\alpha_P = {}^1\!/_2$. Both r_0 and r_1 are monotonically increasing in α_P. Accordingly, when $\alpha_P < {}^1\!/_2$ (the policy-maker is y-leaning), the Lebesgue measure of the set of ideal points of x-leaning experts that the policy-maker would remove is larger than that of the set of ideal points of y-leaning experts that he or she would remove and when $\alpha_P > {}^1\!/_2$ (the policy-maker is x-leaning), the Lebesgue measure of the set of ideal points of y-leaning experts that the policy-maker would remove is larger than that of the set of ideal points of x-leaning experts that he or she would remove, as was to be shown. ∎

The next proposition describes a comparative static of $R(\alpha_P, p, \delta, \kappa, \rho)$ with respect to α_P. While the result is stated in terms of $\alpha_i > {}^1\!/_2$, the result is symmetric once relabeled for $\alpha_P < {}^1\!/_2$.

Proposition 4 *Suppose that $\alpha_{P'} > \alpha_P > {}^1\!/_2$. Then for any $(p, \delta, \kappa, \rho)$, $R(\alpha_{P'}, p, \delta, \kappa, \rho)$ is unambiguously "to the right of" $R(\alpha_P, p, \delta, \kappa, \rho)$:*

$$\min\left[\left\{\alpha_E \in R(\alpha_P, p, \delta, \kappa, \rho)\right\}\right] < \min\left[\left\{\alpha_E \in R(\alpha_{P'}, p, \delta, \kappa, \rho)\right\}\right] \quad \text{and}$$
$$\max\left[\left\{\alpha_E \in R(\alpha_P, p, \delta, \kappa, \rho)\right\}\right] < \max\left[\left\{\alpha_E \in R(\alpha_{P'}, p, \delta, \kappa, \rho)\right\}\right].$$

Proof: Follows from the fact that

$$\frac{\partial x^{\alpha}(1-x)^{1-\alpha}}{\partial \alpha} < 0 \quad \Leftrightarrow \quad x < \alpha, \text{ and}$$
$$\frac{\partial x^{\alpha}(1-x)^{1-\alpha}}{\partial \alpha} > 0 \quad \Leftrightarrow \quad x > \alpha.$$

∎

C Collegial Policy-making

The previous analysis considered only a unilateral policy-maker. We now consider a more general collegial setting in which there are n policy-makers, $N = \{1, 2, \ldots, n\}$, who choose policy as follows.

1. The status quo policy is exogenous and (for simplicity) located at α_E.[27]

2. The policy preferences of the n policy-makers are realized and made common knowledge. We assume that these ideal points are independently and identically distributed according to the cumulative distribution function, F, as defined in Section A.

3. The policy-maker whose ideal point is equal to the median of the policy-makers' ideal points[28] offers a take-it-or-leave-it proposal, z, which is made common knowledge.

4. The policy-makers vote simultaneously with each policy-maker $i \in N$ either voting for ($v_i = 1$) or against ($v_i = 0$) the proposal z.

5. The proposal z is implemented if it receives at least $\tau \geq \frac{n+1}{2}$ votes (*i.e.*, $\sum v_i \geq \tau$). Otherwise, α_E remains in force.

[27]We could have an arbitrary status quo policy, in which case the analysis will of course depend on the location of this policy. This represents an interesting direction (for example, there will be non-obvious situations in which there exists unanimous support for policy movement within the Pareto set) but is beyond the scope of this article.

[28]If there are multiple such policy-makers (either by chance or because n is even, then one is chosen fairly and randomly.

6. The disaster, d, is realized, the players receive their payoffs, and the game ends.

As our solution concept, in order to avoid a multiplicity problem and eliminate a wide range of pathological equilibria, we consider only subgame perfect Nash equilibria in which all policy-makers use weakly undominated voting strategies (**?**, Patty, Snyder and Ting (2009)).

Majority Rule ($\tau = \frac{n+1}{2}$). Individual induced preferences over the set of Pareto efficient policies for any budget B are single-peaked in this setting and thus admit a "representative policy-maker." Specifically, the preferences of a majority of the policy-makers will be identical to those of the policy-maker with the median value of α_i. Accordingly, a key point at the beginning of any analysis of collegial decision-making is that the distribution of the median policy-maker, under the assumption that the policy preferences of the collegial policy-makers are independently distributed according to the cumulative distribution function, F, as defined in Section A, is increasingly less polarized as the number of the number of policy-makers, n, increases. This leads to the following comparative statics.

Proposition 5 *Suppose that the expected location of the median of the policy-makers' preferences (*i.e., $\alpha_m(F) \equiv \{x : F(x) = {}^1\!/_2\}$*) is equal to* ${}^1\!/_2$. *Then for any setting* $(p, \delta, \kappa, \alpha_E, \rho)$*, the probability of replacing an expert, given his or her policy preference,* $\alpha_E \in (0,1)$*, is decreasing in* n *for* n *sufficiently large if*

$$EU(z_P^R(\alpha_E, \rho) \mid \alpha_P, 1) > EU({}^1\!/_2 \mid {}^1\!/_2, p, \delta, \kappa, 0)$$

and increasing in n *for* n *sufficiently large if this inequality is reversed.*

Proof: The result follows from the fact that the independence of the policy-makers' ideal points implies that, for any $\varepsilon > 0$, the probability that the median of the policy-makers' ideal points is located outside of $({}^1\!/_2 - \varepsilon, {}^1\!/_2 + \varepsilon)$ goes to zero as $n \to \infty$. ∎

Proposition 5 can be extended in intuitive ways to allow for the expected value of the median of the policy-makers' preferences to be something other than ${}^1\!/_2$, of course.

Supermajority Rule ($\tau > \frac{n+1}{2}$). If $\tau > \frac{n+1}{2}$, then in line with the "pivotal politics" analysis in Krehbiel (1998), a more supermajoritarian (*i.e.*, less decisive) rule for collegial decision-making will result in more deference to the experts.[29]

Proposition 6 *For any $\alpha_E \in (0,1)$, the probability of replacing α_E is decreasing in the threshold τ.*

Proof: Fix $(p, \delta, \kappa, \rho, F)$. In equilibrium, α_E will be replaced if and only if the profile of policy-maker ideal points, $\alpha_N \equiv \{\alpha_!, \dots, \alpha_n\}$ satisfies both of the following:

$$\begin{aligned} \text{median}[\alpha_N] &\in \{i \in N : \alpha_i \in M(p,\delta,\kappa,\alpha_E,\rho)\}, \text{ and} \\ \tau &\leq \Big|\{i \in N : \alpha_i \in M(p,\delta,\kappa,\alpha_E,\rho)\}\Big|. \end{aligned} \tag{2}$$

The probability that condition (2) is satisfied is

(3)

$\sum_{t=\tau}^{n} \binom{n}{t} F(M(p,\delta,\kappa,\alpha_E,\rho))^t (1 - F(M(p,\delta,\kappa,\alpha_E,\rho)))^{n-t}$.By the assumption that F assigns positive measure to all open sets in $[0,1]$ and the fact that $M(p,\delta,\kappa,\alpha_E,\rho)$ is a nonempty interval strictly contained in $[0,1]$ (Proposition 1), it follows that $F(M(p,\delta,\kappa,\alpha_E,\rho)) \in (0,1)$, so that (3) is strictly decreasing in τ, as was to be shown. ∎

[29] It is important to note that the ease of implication relies on the fact that individual preferences over the set of Pareto efficient policies are single-peaked. On this point and thoughts about extending Krehbiel's analysis to include policy-specific "valence" characteristics (which the retention of expertise in our framework is a version of), see Monroe, Patty and Penn (2018).

D Uncertainty About the Value of Expertise

Information Structure. Suppose now that the policy-maker is uncertain about κ. Specifically, suppose that the true value of κ is drawn from a distribution of the following form:

$$\kappa = \begin{cases} k_1 & \text{with probability } \phi, \\ k_0 & \text{with probability } 1 - \phi, \end{cases}$$

where $0 \leq k_0 < k_1 \leq 1$ and $\phi \in (0, 1)$ are exogenous and common knowledge. Furthermore, prior to the decision sequence described in the baseline model in Section 3, the policy-maker observes a signal, $\sigma \in \{k_0, k_1\}$, realized according to the following distribution conditional on κ:

$$\Pr[\sigma = \kappa \mid \kappa] = q,$$

where $q > 1/2$ is exogenous and common knowledge.

Interpreting σ and q. As we return to below, the parameter q represents the degree to which the policy-maker can "trust" the information contained in the signal, σ. When $q \approx 1/2$, the signal contains little information, whereas when $q = 1$, the signal is *known* to be true. In the context of this article, σ might represent reports about the capacity of the agency (or agencies) responsible for mitigating the effects of a disaster. In this interpretation, $\sigma = k_0$ represents a report that suggests that the agency has low capacity, and $\sigma = k_1$ represents a report suggesting that the agency possesses high capacity.

In this interpretation, then, q represents the degree to which the policy-maker changes his or her beliefs about the agency's capacity after observing the report, σ. Thus, a policy-maker who "doesn't trust" the source of σ would be represented by presuming that q is close to $1/2$.[30]

[30]For reasons of space, we are sidestepping a subtle epistemological question here. Specifically, the traditional non-cooperative game theoretic approach would assume that q *is* the true reliability of σ. In the unilateral policy-making setting considered in Section 4, this assumption is required only insofar as one wants

Equilibrium Analysis. As in the baseline model in Section 3, the policy-maker's optimal choice must be either z_P^R or z_1^*. The question of which one is optimal can depend on the policy-maker's posterior belief about κ conditional of his or her signal, σ. Thus, in this setting, standard arguments (*e.g.*, Austen-Smith and Banks (1996)) imply that, in equilibrium, the unilateral policy-maker will respond to the information contained in his or her signal, σ, only if q is sufficiently high.[31] Otherwise, he or she will act essentially based only on his or her prior belief, ϕ.

As earlier, we will denote a situation in this extended model by $\xi \equiv (\alpha_P, \alpha_e, p, \delta, \kappa, e, \sigma, \phi, q)$.

Denoting the posterior belief that $\kappa = k_1$, conditional on σ, ϕ, and q, by

$$\beta(\sigma) \equiv \beta(\sigma \mid \phi, q) = \begin{cases} \frac{\phi q}{\phi q + (1-\phi)(1-q)} & \text{if } \sigma = k_1, \\ \frac{\phi(1-q)}{\phi(1-q)+(1-\phi)q} & \text{if } \sigma = k_0, \end{cases}$$

the expected location of the policy outcome in the absence of expertise, given the policy choice z and signal σ is[32]

$$\bar{z}(z, e, \sigma \mid \xi) = (1 - \delta(1 - ((k_1 - k_0)\beta(\sigma) + k_0)e))z.$$

With this in hand, the policy-maker's expected payoff from a policy z in a situation ξ (presuming that $e = 0$) is

$$\begin{aligned} \overline{EU}_1(z \mid \sigma, \xi) &= (1-p)u_1(z) + p \cdot u_1(\bar{z}(z, 0, \sigma \mid \xi)), \\ &= (1 - p\delta)u_1(z), \end{aligned}$$

to consider welfare questions—the unilateral policy-maker can never be better off having an incorrect value of q. In richer settings, however, the assumption can be incredibly important.

[31] A key commonality between this setting and that of Austen-Smith and Banks (1996) is that the set of choices that can be sequentially rational for the policy-maker is binary.

[32] Recall that the Cobb-Douglas payoff function implies that the policy-maker is risk-neutral with respect to the budget constraint, so that the expected policy payoff equals the payoff from the expected policy outcome.

so that, denoting the expected level of κ, conditional on σ, by

$$\bar{\kappa}(\sigma) = \beta(\sigma)(k_1 - k_0) + k_0,$$

the policy-maker should choose $z = z_P^R$ after signal σ if

$$u_1((1-p(1-\bar{\kappa}(\sigma))\delta))z_P^R) \quad > \quad (1-p\delta)u_1(z_1^*).$$

Thus, the policy-maker should be responsive to his or her signal, σ, if the following conditions simultaneously hold:

$$\frac{1-p+p\bar{\kappa}(0)\delta}{1-p\delta} < \frac{u_1(z^*)}{u_1(x_P^R)} < \frac{1-p+p\bar{\kappa}(1)\delta}{1-p\delta} \tag{4}$$

With this in hand, the key points of this appendix are stated formally in the following proposition.

Proposition 7 *Policy will be responsive to the perceived level of effectiveness of expertise in equilibrium only if*

1. z_P^R *is sufficiently far from* z_P^*,
2. $q \in (1/2, 1]$ *is sufficiently large,*
3. $k_1 - k_0 > 0$ *is sufficiently large, and*
4. $p \cdot \delta$ *is neither too small nor too large.*

Proof: Each of the conclusions follows directly from (4). ∎

The first conclusion of Proposition 7 is simple, but important: deference to the expert will be unconditional (*i.e.*, independent of the policy-maker's beliefs about the value of the expert's expertise) so long as the expert's policy preferences are similar enough to those of the policy-maker.

The second conclusion of Proposition 7 implies that the policy-maker will be responsive to his or her information about the value of expertise only if that information is sufficiently reliable. For reasons of space, we do not explore the foundations of this reliability (*i.e.*, q), but it is worth noting that the implication is in line with the policy-maker ignoring advice from sources about the value of expertise from a source that he or she does not "trust." In addition, it provides foundations, beyond direct ideological motivation, for inconsistent patterns in turnover across Administrations. Some new presidents might find it harder to trust existing experts, either for reasons of ideological preferences or for reasons related to expert experience, personal connections, or familiarity with the expert's background. Moreover, to the extent that a politician might have a view of q that is affected by past policies, outcomes, or decisions an expert has made, we might expect to observe a connection between expertise turnover and the salience of substantive policies.

The third conclusion of Proposition 7 is intuitive—the policy-maker will be responsive to his or her signal only if the stakes are sufficiently large. An interesting ancillary implication of this is that the minimal level of these stakes (*i.e.*, the minimal level of $k_1 - k_0$) required for responsiveness is increasing in the divergence between the policy preferences of the expert and those of the policy-maker and decreasing in the size of an unmitigated disaster (δ).

Finally, the fourth conclusion of Proposition 7 is based simply on the fact that when $p\delta \approx 0$, the policy-maker should choose z^* regardless of the value of κ and when $p\delta \approx 1$, he or she should choose z_P^R regardless of the value of κ.

E The Expert's Career and Policy Motivations

The analysis in the body of the article treats the expert E's motivations in a reduced-form version—there is an exogenous reservation level, ρ, that the expert can obtain by leaving the government. We present a fuller model of this reservation level that incorporates the expert's policy preferences outside government.[33]

[33] We thank an anonymous reviewer for suggesting this extension.

Suppose that E is paid a governmental wage of $g > 0$ and, if E leaves government service, E can secure an (expected) outside wage denoted by $w > 0$. Then suppose that E's overall payoff function is[34]

$$\pi_E(e \mid z, g, w, p, \delta, \kappa, \alpha_E) = e \cdot g + (1 - e) \cdot w + (1 + \zeta e) EU(z \mid \alpha_E, p, \delta, \kappa, e),$$

where $\zeta \geq 0$ measure the degree to which E cares "more about" the policy impacts of z when he or she is responsible for implementing z.[35]

The expert's optimal response, $e^*(z \mid g, w, p, \delta, \kappa, \alpha_E) \in \{0, 1\}$ is essentially determined by the following rule:

$$e^*(z \mid g, w, p, \delta, \kappa, \alpha_E) = \begin{cases} 1 & \text{if } EU(z \mid \alpha_E, p, \delta, \kappa, 1) \geq \frac{w-g}{\zeta}, \\ 0 & \text{if } EU(z \mid \alpha_E, p, \delta, \kappa, 1) < \frac{w-g}{\zeta}. \end{cases}$$

Now, letting

$$\rho \equiv \frac{w - g}{\zeta},$$

one can include both policy preferences and work preferences for E within the reduced form model in the body of the article. The key point, of course, is to demonstrate that explicitly incorporating policy preferences on the part of E need not imply that E will always remain in government service.

Comparative Statics. We can say more at this point about the effects of both the new parameters (g, w, and ζ) as well as those in the model as presented in the body of the

[34]We include the full set of parameters (specifically, p, δ, κ) here because we will discuss the "comparative statics of ρ" with respect to each. Also, we do not need to assume that E's payoffs are quasi-linear with respect to the wage—any increasing function of the wage would work just as well, at the expense of greater notational bloat.

[35]Note that, unlike with typical quadratic loss policy preferences, $EU(z \mid \alpha_E, p, \delta, \kappa, e) \geq 0$, so that $\zeta > 0$ implies that it is always preferable to work for the government if $g = w$. We think that this formulation appropriately captures the notion that individuals get consumption value from carrying out a job that is more in line with their own values.

article (p, δ, κ, and α_E). To make this precise, suppose that the expert privately observes an idiosyncratic private sector wage,

$$w = W + \varepsilon,$$

where $W > 0$ and $\varepsilon \sim \text{Normal}(0, s^2)$ for some exogenous standard deviation $s > 0$. Then

$$\Pr[e = 1] = \Pr\left[\varepsilon \leq \zeta \cdot EU(z \mid \alpha_E, p, \delta, \kappa, 1) + g - W\right]. \tag{5}$$

The following proposition follows immediately from (5).

Proposition 8 *For any fixed policy z, the probability that the expert will be retained in government service is*

1. *decreasing in p,*
2. *decreasing in δ,*
3. *increasing in κ,*
4. *decreasing in the distance between α_E and z,*[36]
5. *increasing in g,*
6. *decreasing in W, and*
7. *increasing in ζ.*

F Incorporating Electoral Accountability

The model in the body of the article elides the issue of electoral accountability—the distribution of α_P is treated as both the distribution of the median voter's ideal point and the ideal point of the policy-motivated executive. There are many ways one might expand the

[36]This is a bit loose. Formally, it is decreasing in the distance "on either side" of α_E. Unless $\alpha_E = (1/2, 1/2)$, the actual decrease for a given positive distance between z and α_E will depend on which side of α_E z is on.

model to consider the challenges of electoral accountability and their effects on deference to expertise in more detail. We consider only a very simple one in this appendix.[37]

Specifically, extend the model to two periods and now include a third player, a voter V, with (privately observed) ideal point $\alpha_V \in [0,1]$, which is drawn from the same distribution, F, as defined in Appendix A. The role of the voter is to choose, after observation of z, e, and d, whether to reelect the policy-maker ($r = 1$) or not ($r = 0$). If V chooses $r = 1$, then (for simplicity) he or she receives his or her expected policy payoff from z again, conditional on e.[38] If V chooses $r = 0$, on the other hand, a new policy-maker (the "challenger," C) is elected, with ideal point α_C drawn from the same distribution as α_V. In addition, the government's expertise is retained/restored, so that the challenger can set policy so as to have $e = 1$ in the second period. The first-period (or, "incumbent") policy-maker's preferences are the same as before, but with some (potentially type-specific) value for reelection, $W(\alpha_P) > 0$.[39]

In such a model, the voter will vote for the incumbent policy-maker if and only if the voter's expected payoff from z and e is high enough. (The model is "hard-wired" to make "retrospective voting" the best response for V.) Because $e = 1$ Pareto dominates $e = 0$ for any fixed z, this will induce the incumbent to be less willing to choose a policy z that induces E to depart from government service. However, so long as $\min[\lim_{\alpha_P \to 0} W(\alpha_P), \lim_{\alpha_P \to 1} W(\alpha_P)]$ is not sufficiently large, there will exist a positive measure of incumbent ideal points that will set policy $z = z^*(\alpha_P)$ and induce E to leave government service.

Furthermore, if $W(\alpha_P) = \bar{W} + \bar{k} EU(z_1^*(0))$ for $\bar{W}$ sufficiently small and $\bar{k} < 1$, all of the comparative statics identified in the body of the article will remain true in this extended setting.[40]

[37]We thank an anonymous reviewer for suggesting this extension.

[38]That is, nothing changes in the second period.

[39]Note that this is distinct from the parameter W in Appendix E. Also, note that allowing $W(\alpha_P)$ to depend on the α_P implies that the results are consistent with purely policy-motivated policy-makers in addition to possibly capturing "ego rents" from reelection.

[40]This is not a necessary condition—the key point is that the behavior of $W(\alpha_P)$ with respect to α_P can not deviate "too much" from that of EU in order for the comparative statics of the baseline model to be preserved.

G Two Unidimensional Spatial Examples

Suppose that the policy space is a subset of $\mathbf{R}$ and each player i has a type $\alpha_i \in \mathbf{R}$ and policy preferences are *spatial* in the sense that, for any $(d, e) \in \{0, 1\}^2$,

$$|z - \alpha_i| < |z' - \alpha_i| \Leftrightarrow u_i(z \mid \alpha_i, d, e) > u_i(z' \mid \alpha_i, d, e).$$

There are at least two ways to incorporate the value of expertise in such a setting. The first route is to treat expertise as a "valence" good that is valued by all players, *ceteris paribus*. The second way is to presume that the set of feasible policies following a disaster depends on whether the expert is present or not.

The causal theory offered in the body of the article is that experts help disaster response by enlarging the set of feasible policies following a disaster. In the policy environment considered in the body of the article, it is always Pareto efficient to enlarge this set (in the sense of set inclusion). This assumption is made to clarify the strength of the argument: we demonstrate that some policy-makers voluntarily induce the expert to leave the government in pursuit of policy, in spite of the fact that the departure of the expert may ultimately hurt the policy-maker's policy goals. Assuming that the expert simply represents a valence good retains the Pareto efficiency of retaining the expert, but divorces the origin of this efficiency from the policy-maker's policy-seeking motivations. Furthermore, it leaves the question of the origins of the valence characteristic of expertise unmotivated. Indeed, such a model (as presented below in Example 9) is equivalent to there being a positive fixed cost of moving policy (*i.e.*, transaction costs/"policy stickiness").

Example 9 (Expertise as a Valence Good) Suppose that each player has the following payoff function:

$$u_i(z \mid \alpha_i, d, e) = (e\kappa - \delta)d - |z - \alpha_i|,$$

where $\kappa > 0$ measures the value of expertise in a disaster. In such a model, setting policy farther than ρ from α_E will induce the expert to leave government and impose an expected cost of $p\kappa$ on the policy-maker. Accordingly, in equilibrium, the policy-maker will be willing to induce the expert to leave government only if

$$|\alpha_P - \alpha_E| \geq \rho,$$

and (supposing without loss of generality that $\alpha_E + \rho < \alpha_P$)

$$\alpha_P - \alpha_E - \rho \geq p\kappa.$$

It is simple to see that this model will generate all of the key conclusions of the model presented in the body of the article (*i.e.*, Predictions 1, 2, and 3). △

The next example illustrates that having the expert enlarge the set of feasible policies after a disaster will not yield the same predictions as the model presented in the body of the article.

Example 10 Let $Z(d,e) \subseteq \mathbf{R}$ denote the feasible set of policies given disaster occurrence d and expertise level e and suppose that each $Z(d,e)$ is a closed interval, satisfying the following:

$$Z(1,0) \subset Z(1,1) \subset Z(0,0) \subseteq Z(0,1) = \mathbf{R}, \tag{6}$$

and let $\mathcal{Z} \equiv \{Z(0,0), Z(0,1), Z(1,0), Z(1,1)\}$ denote the set of sets of feasible policies. The ordering in (6) then captures the idea that disasters restrict the set of feasible policies, and expertise mitigates the shrinkage induced by a disaster. The remainder of the model is identical to the model presented in the body of the article. The analogue to $M(p,\delta,\kappa,\alpha_E,\rho)$ from the original model in this setting (the set of policymaker's ideal points that would induce the policy-maker to defer to an expert with ideal point α_E) is the following:

$$M_{\text{Spatial}}(p,\delta,\alpha_E,\rho,\mathcal{Z}) = \left\{\alpha_P \in \mathbf{R} : EU(z_P^R(\alpha_E,\rho) \mid \alpha_P, 1) > EU(z_1^*(0) \mid \alpha_P, p, \delta)\right\}.$$

In this setting, we can define "moderate" preferences with respect to $\mathcal{Z}$ as follows: the most "moderate" policies are those that are always feasible: policy $z \in \mathbf{R}$ is moderate if $z \in Z(1,0)$. The following proposition implies that, in the unidimensional spatial setting, "moderate" policymakers will retain "moderate" experts if and only if they have *identical* policy preferences.

Proposition 11 *In the unidimensional spatial setting, for any* $(p, \delta, \rho, \mathcal{Z})$ *and any* $\alpha_E \in \text{Int}(Z(1,0))$,

$$\left(\alpha_P \in Z(1,0) \cap M_{\text{Spatial}}(p, \delta, \alpha_E, \rho, \mathcal{Z})\right) \Leftrightarrow \left(\alpha_P = \alpha_E\right).$$

Proof: Fix $(p, \delta, \rho, \mathcal{Z})$ and any $\alpha_E \in \text{Int}(Z(1,0))$, and consider any $\alpha_P \in (Z(1,0))$. The policy-maker's payoff from choosing $z = \alpha_P$ achieves its maximum with probability 1, because a disaster will not eliminate P's ideal point. ∎

Proposition 11 implies that policy-makers with moderate preferences will be no more likely (and frequently strictly less likely) than those with extreme preferences to defer to expertise, at odds with Prediction 1 of the model in the body of the article. △