2011

***Left, Right, and Center.* Strategic Information Acquisition and Diversity in Judicial Panels**

Matthew L. Spitzer
mspitzer@law.usc.edu

Eric L. Talley
Columbia Law School, etalley@law.columbia.edu

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship

Part of the [Judges Commons](#)

Recommended Citation

Matthew L. Spitzer & Eric L. Talley, *Left, Right, and Center: Strategic Information Acquisition and Diversity in Judicial Panels*, JOURNAL OF LAW, ECONOMICS & ORGANIZATION, VOL. 29, P. 638, 2013; UC BERKELEY PUBLIC LAW RESEARCH PAPER NO. 1832765 (2011).

Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1693

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact scholarshiparchive@law.columbia.edu.

*Left, Right, and Center: Strategic Information Acquisition and Diversity in Judicial Panels**

Matthew Spitzer[†] & Eric Talley[‡]

June 25, 2011

Abstract

This paper develops and analyzes a hierarchical model of judicial review in multimember appellate courts. In our model, judicial panels acquire information endogenously, through the efforts of individual panelists, acting strategically. The resulting equilibria strongly resemble the empirical phenomena collectively known as “panel effects” – and in particular the observed regularity with which ideological diversity on a panel predicts greater moderation in voting behavior (even after controlling for the median voter’s preferences). In our model, non-pivotal panel members with ideologies far from the median have the greatest incentive to acquire additional policy-relevant information where no one on a unified panel would be willing to do so. The resulting information structure pushes deliberation and observed voting patterns towards apparent moderation. We illustrate the plausibility of our model by calibrating it to empirical data, and explore various normative implications of our theory.

*(First Version: May 6, 2009; Submitted for Publication: Aug. 23, 2010). Thanks to Bernard Black, Charles Cameron, Steve Choi, Robert Cooter, Lee Epstein, John de Figueiredo, Rui de Figueiredo, Dan Farber, Joshua Gans, Jake Gersen, Al Klevorick, Prasad Krishnamurthy, Tracy Lewis, Justin McCrary, John Morgan, Anne Joseph O’Connell, Kevin Quinn, Mark Ramseyer, Suzanne Scotchmer, Pablo Spiller, Matthew Stephenson, Cass Sunstein, Emerson Tiller, Abe Wickelgren, Jonathan Wiener, seminar participants at the American Law and Economics Association meetings, Berkeley, Chicago, Northwestern, Melbourne Business School, Texas, Vanderbilt, and the Duke-North Carolina Triangle Law and Economics Conference for comments and discussions. We also thank two anonymous referees for constructive suggestions. Sam Houshower, Teddy Maghee, and Daniel Mitchell provided excellent research assistance on this project. All errors are ours.

[†]University of Texas. Email: mspitzer@law.utexas.edu

[‡]University of California, Berkeley. Email: etalley@law.berkeley.edu.

1 Introduction

Within the growing empirical literature on judicial review, three notable findings stand out. First, *politics matters*: Judges appointed by Democrats are more likely to uphold liberal agency and/or trial court decisions and reverse conservative ones than are their Republican-appointed counterparts.¹ Second, *party matters*: while manifesting qualitatively similar behavior, Democrat and Republican judges do not mirror one another exactly (*e.g.*, they do not always “cross the party line” with the same frequency). Third, *diversity matters*: mixed three-judge panels (*i.e.*, two Democrats and one Republican or two Republicans and one Democrat) tend to make decisions that are more moderate than do homogenous panels dominated by a single party (Democrat or Republican).²

This paper focuses on the third feature – the evident moderating effects of panel diversity – and in the process says something about the other two. Our contribution is primarily theoretical: we develop and analyze a model connecting (a) hierarchical auditing of lower-tier actors (*e.g.*, administrative agencies or trial courts); (b) group deliberation within a multimember auditing entity (*e.g.*, an appellate judicial panel); and (c) strategic decisions by group members to make costly investments in information acquisition relevant to deliberation (*e.g.*, about the case itself, underlying policy choices at play, doctrinal constraints, etc.). Our model predicts each of the empirical regularities noted above as an equilibrium phenomenon, and in particular the evident moderating effect of panel diversity. Specifically, we show that ideologically heterogeneous panels are more likely to incentivize broad information production than are homogenous ones. For example, a lone Republican (or Democrat) on a three-judge panel may be willing to uncover new information for her counterparts even if they are not willing to provide it for themselves. The information flow induced by panel diversity generates voting patterns that appear more moderate than those of homogeneous panels. To the extent that our hypothesis is correct, it holds implications for whether mixed judicial panels should be encouraged and/or required. (Miles and Sunstein 2008, Tiller and Cross 1999; cf. Schanzenbach and Tiller 2008, 2009).³

¹See Revesz (1997), Cross and Tiller (1998), and Miles and Sunstein (2006, 2008), Sunstein et al (2004), as well as earlier work in political science (such as Songer 1982), for empirical confirmation. The explanation for this phenomenon is widely thought to be ideological disposition. Segal and Spaeth 2002. Accord Stephenson (2009) at 46.

²See also Peresie (2005), finding similar effects for male and female judges. For an excellent overview, Hettinger et al. (2007). For the history of this phenomenon, see Maveety (2005) and Kastlelec (2008). Kastlelec produces data suggesting that panel effects are a comparatively recent phenomenon, arising in the second half of the 20th Century. By 2011, however, there is no doubt that robust panel effects exist.

³As Stephenson (2009: pg. 47) points out, there are two effects from mixed judicial panels. One is the tendency of the minority judge to vote with the majority. The second, and in our opinion likely the more important effect, is the tendency of the majority judges to creep ever

The framework we develop builds on our prior work (Spitzer and Talley 2000),⁴ but departs from it in a few crucial ways. First, we generalize the model to yield a deeper understanding of appellate court dynamics (and hierarchical auditing more broadly). Rather than treating the appellate court as a unitary actor (as both we and Cameron et al. (2000) did), we explicitly model it as a multimember body. This generalization is critical, since strategic interaction among panelists generates the core intuitions we highlight here. Second, we tailor our framework to include key institutional attributes in administrative law. When agency / trial court decisions are appealed, the court must hear such appeals. Yet, appellate judges can choose how intensively to scrutinize the matter before them. Our model specifically captures this endogenous effort choice by individual judges sitting on a larger panel. Finally, our framework can be simulated and tested with real-world data, and accordingly we demonstrate that a calibrated version of our model predicts patterns of panel effects that are observed in the existing literature.

Political scientists have suggested a number of theories for explaining the moderating effect of ideological diversity on three judge panels. A first set of explanations hinges on *social cohesion and collegiality* (e.g., Songer 1982: 226), positing that social pressures may lead non-pivotal minority judges to go along with the majority, as a mechanism for enhancing (or preserving) inter-panelist harmony. Even if such tastes for collegiality are relatively weak, they may be enough to deter the minority panelist from taking the time and energy to author a dissent.⁵ *Dissent aversion* – a set of predictions about when judges will allocate their time to writing dissents – partially relies on a theory of social cohesion and collegiality. Epstein et al. (2011), for example, show that dissent incidence is negatively associated with caseload and positively associated with both circuit size and intra-circuit ideological diversity, all of which may bear on the costs, benefits and sustainability of collegial norms among appellate court judges. In a related vein, some have posited that additional pressures from *group polarization* may play a more extreme role in homogenous panels, which can in turn lead to apparent moderation of mixed panels (e.g., Sunstein et al., 2004: 308). That is, individuals may become more extreme when interacting with like minded counterparts (Myers 1975; Asch 1951). Applied to judges, polarization effects predict that homogenous panels reinforce each other's prior

so subtly in the direction of the minority. As it happens, in our model, described below, both effects can occur simultaneously.

⁴We are also implicitly building on Cameron et al. (2000), which was published contemporaneously with Spitzer and Talley (2000), and uses a model very similar in spirit.

⁵Within this literature, both social and workload-related costs/benefits can play a role. Atkins (1973) ("social pressure"); Atkins & Green (1976) (empirical support for workload and dissents inversely related); Goldman (1968) (norm of consensus); Green (1986) (workload reduces dissents). See also Posner (2002: 32) ("[m]ost judges do not like to dissent....Not only is it a bother and frays collegiality, and usually has no effect on the law, but it also tends to magnify the significance of the majority opinion."); Landes and Posner (2009) (discussing "dissent aversion"). Relatedly, Fischman (2009) studies an attitudinal model, augmented by a cost to writing a dissent. The greater that cost, the more likely it is that a minority judge will join the majority opinion. His model does not predict that the majority judges will moderate their position and join the minority.

commitments, thereby leading to more ideologically extreme decision making (and apparently more moderation in mixed panels).

A second explanation, *whistleblowing*, is perhaps the leading explanation among positive political theory (PPT) scholars to explain panel effects. First proposed by Cross and Tiller (1998), this account conjectures that a minority party panelist can effectively threaten to tattle on the majority (e.g., through a dissent) if those majority actors ignore or misconstrue established precedent or doctrine. If the minority member can credibly threaten can expose a majority’s omissions or misstatements, she can deter such behavior in the first instance, producing more moderation. (Cross and Tiller 1998: 2156). The whistleblower account harbors a distinct role for formal legal doctrine as a constraint on judicial review. That is, the whistleblower account gets its traction from the existence of an independent, commonly subscribed legal canon, whose violation can be detected and communicated to a higher authority. Our approach, in contrast, neither requires nor precludes the possibility that legal doctrine might also do some work and, in fact, allows for doctrine to be vague, contested, over- or under-determined, or simply unintelligible. In order to highlight the role of endogenous information production, we will focus only on ideology, information, choice, and outcomes.⁶

A final explanation, perhaps the leading one among legal academics, was proposed by Revesz (1997: 1732), and is sometimes identified as the *deliberation* hypothesis. In essence, by being empaneled with judges from the opposite political party and deliberating with them, one is naturally led to moderate her positions. The informational explanation that we propose here is perhaps closest in spirit to Revesz’ suggestion, but we develop it within a more formal theoretical framework, generating in turn more precise predictions about the mechanics of panel effects.

Before proceeding, one caveat deserves specific mention. Although our analysis aims to understand and explain judicial panel effects, it has obvious ties to other literatures in political science, psychology, economics and elsewhere on group deliberation. These include papers on (so called) persuasion games,⁷

⁶We hasten to add that the role of doctrine may certainly be important (and we have published on the role and characterization of doctrine before (Cohen & Spitzer (1994); Talley (1999)). Indeed, the process of “information acquisition” we model here could, in principle, also extend to conventional legal research on existing precedent, so that whistleblowing would represent a special case within our framework. We take up this interpretation below.

More generally, the strategic formulation of doctrine by the Supreme Court, and its effects on lower courts, takes up a significant fraction of research in this field. See Cross & Tiller (2006), Jacobi (2009), Jacobi and Tiller (2007), Lax (2007, 2008), Lax and Landa (2009), McNollGast (1995), Strauss (1987), Tiller (1998), Shavell (2009), and Stephenson (2006). Rodriguez and Weingast (2007) have extended this approach to the interaction between the courts and Congress. Kestellec (2007) extends Cross and Tiller’s whistleblower model into Kornhauser’s (1992a, b) “case space” and explores how three-judge panels (as opposed to individual judges) enhance the Supreme Court’s ability to see its preferred doctrine carried out. Kestellec (2010) extends the model to a two-level hierarchy above three judge panels, and finds the asymmetric form of control induced by whistleblowing that Cross and Tiller (1998) predicted.

⁷Milgrom and Roberts (1986).

inquisitorial versus advocacy systems,⁸ political lobbying,⁹ media reporting and bias,¹⁰ and the value of ideological diversity more generally within deliberative fora.¹¹ We do not attempt to develop these links fully here, though our general approach may both inform such inquiries and is, in many respects, informed by them.

Our analysis proceeds as follows. Section 2 describes at greater length the literature relating to panel effects, along with the prevailing theories that have been posited to explain them. Section 3 presents our theoretical model and characterizes its equilibria. Section 4 uses simulation methods to calibrate our model to existing empirical data, and develops some preliminary thoughts about testing our model against alternatives. Section 5 discusses extensions of our model. Section 6 considers implications, and Section 7 concludes.¹²

2 Empirical Panel Effects

Before beginning with our analytic enterprise, it is perhaps useful to situate our claims within the empirical literature. As noted in the introduction, during the last decade the empirical literature on judicial panel effects has proliferated rapidly. Although we cannot review this field in its entirety, a few of the central landmarks in this literature are worth recounting. Revesz (1997) is often credited with being the first legal academic to notice and document the phenomenon. He collected challenges to decisions of the Environmental Protection Agency that were brought in the DC Circuit between 1970 and 1994. Revesz divided the time into periods in which the membership of the DC Circuit was unchanged and utilized the random assignment of judges to test hypotheses about the effect of panel composition on votes and outcomes.¹³ Revesz’s results provided early empirical support for panel effects, though they are mixed as to which pattern of effects is supported by the data. In the 1970s, the link between panel diversity and moderation was manifest in Republican-appointed judges, but not Democrats. In later periods, however, both Republicans and Democrats appeared to display such patterns.

Shortly after Revesz’s study, Cross and Tiller (1998) analyzed 170 cases in which the DC Circuit reviewed agency interpretations of regulatory statutes. They found that unified panels (consisting of all Republicans or all Democrats) were 17% less likely to defer to agencies than were split panels. More pertinent to the panel effects literature, they also found that unified panels deferred to agencies only 33% of the time when the panel’s and agency’s politics were inconsistent, but deferred to the agency 62% of the time when the panel was

⁸Dewatripoint and Tirole (1999).

⁹De Figueiredo and Cameron (2008).

¹⁰Gentzkow and Shapiro (2006).

¹¹Farhang and Wawro (2004); Peresie (2005).

¹²The Appendix includes a number of technical derivations and proofs that are suppressed in the text.

¹³Revesz also tested hypotheses unconnected to panel composition, and found voting patterns that are consistent with an ideological component to judicial voting.

politically heterogeneous (Cross and Tiller, 1998: 2172). This finding is also consistent with the moderating effect of split panels.

Miles and Sunstein (2006, 2008) also present evidence supporting panel effects.¹⁴ They collected all Circuit Court reviews of EPA and NLRB decisions between 1996 and 2006 for insufficient factual basis or for being arbitrary or capricious, which together they call “arbitrariness” review. Next, they compute the *validation rate*, which is the rate at which appellate panels upheld administrative decisions against challenge. Then they coded for the politics of the administrative action by considering who challenged Agency action; if industry challenged the Agency action then the Agency action was deemed liberal, whereas if a union or an environmental group challenged an Agency action, then the Agency action was deemed conservative. Last, Miles and Sunstein coded each judge’s political party as equal to the party of the appointing president for that judge.

Miles and Sunstein found the same basic ideological component of voting that others have found. Judges appointed by Democrats were more likely to validate liberal administrative Agency actions than conservative actions. Judges appointed by Republicans had the reverse tendency. But in addition, Republican appointees were more likely to validate conservative administrative agency actions when they were sitting with two other Republican Judges than when they were sitting with one or more Democrats. Democrat appointees appeared to behave in similar (but perhaps more complicated) ways.¹⁵

Unfortunately, Miles and Sunstein constructed their measures by pooling all mixed panels, rather than separating, for example, RRD and DDR panels. On the basis of their articles, then, one cannot directly observe the change in voting tendencies between a sole minority member of a panel and the same judge who enjoys a two-judge majority. Nevertheless, Miles and Sunstein measure the empirical propensity of a Democrat judge to uphold agency decisions when she is moved from a unified Democrat panel to a mixed panel. They find Democrat appointees are less likely to validate liberal agency decisions – and more likely to uphold conservative decisions – when deliberating in mixed panels. Republican propensities move in the opposite qualitative direction (although Republican voting patterns were somewhat less sensitive to panel composition than were Democrats’). We regard these results as evidence in favor of panel effects; that is, inclusion on a mixed panel tends to moderate voting patterns.¹⁶

¹⁴Because we later utilize the Miles & Sunstein data to calibrate our model, we spend some time here describing their papers in detail.

¹⁵Similarly, Sunstein et al. (2004) investigated federal appeals panels in thirteen categories. They found that the typical pattern of panel effects existed in most of the subject areas (e.g. campaign finance, affirmative action, EPA regulation); however, in at least one context (Title VII discrimination cases) it was muted, and in three areas (federalism, criminal law, takings clause) the pattern was missing entirely. In some of the areas the effects were symmetric, while in other areas not. In two areas (abortion and capital punishment) they found pure ideological voting, but no panel effects at all.

¹⁶It is less clear whether the Miles and Sunstein’s results should be taken as evidence of symmetry or asymmetry between Republicans and Democrats (and could be consistent with either). Cf. Schanzenbach and Tiller (2008).

Landes and Posner (2009) correct and clean the most commonly used data bases, and then present a large number of empirical analyses on judicial review. They claim that they could not code lower Federal Court votes as majority or dissent, and hence they could not say much about panel effects *per se*. They did find, however, that judges appointed by Democrats were more likely to cast liberal votes than were judges appointed by Republicans, and also that mixed panels appeared to manifest “moderation” in their views, at least among Federal Circuit panels (but not on the Supreme Court).

In a significant recent paper that both reviews and contributes to the literature on gender effects in judging, Boyd et al. (2011) tested whether male and female judges vote differently in thirteen different doctrinal areas. They found significant panel effects in only one area: sex discrimination in employment, where males were far more likely to vote liberally when sitting with a female judge than when sitting with only other males. Boyd et al. interpret this result as reflecting an informational explanation of panel effects. Women have information about how employment discrimination works, which they can share with their panelists. Their interpretation meshes very nicely with the mechanism driving our model.¹⁷

Some recent academic contributions have injected some skepticism (or at least words of caution) into the enterprise of empirical estimation of judicial preferences. Edwards and Livermore (2009: 1916), for example, strongly criticize this literature, partly on the ground that it is based on an attitudinal model that does not take into account the dynamics of group deliberation. Our model is the first that we know of to attempt to characterize an important characteristic of deliberation – information exchange. For reasons that are not clear (at least to us), several commentators seem to regard collegial deliberation as inconsistent with ideological explanations. (Edwards and Livermore (2009: 1917); Tacha (1995: 586); Wald (1999: 255)). As our model shows, however, the two concepts not only can coexist, but their interaction may be key to understanding panel effects.

In sum, the empirical literature provides overwhelming support for the proposition that ideological differences among judges matter for outcomes. It also provides significant evidence for a moderation effect in heterogeneous panels, where minority and majority factions tend to move towards one another when voting. Finally, there is some intermittent evidence that even as they exhibit qualitatively similar patterns, Republican and Democrat judges do not always behave as complete mirror images of each other. That said, the precise drivers of these phenomena are still not well understood. Accordingly, the next sections of this paper develop a theoretical framework that can both predict and explain the observed phenomena.

¹⁷Kastellec (2011, working paper) finds similar results for African American Judges on the Courts of Appeal for affirmative action cases.

3 Model

In this section, we introduce and analyze a formal model of deliberation within multijudge panels. Using this model, we show how ideological diversity can endogenously affect information production, which in turn generates panel effects even when judicial preferences remain constant. To expose our key intuitions, we begin with a relatively simple structure, addressing more complicated extensions in later sections.

3.1 Framework

Consider a two-level hierarchy, consisting of a unitary initial actor, A , representing an administrative agency or a trial court, and a multi-member body, J , which represents an appellate judicial panel. We suppose the actors must make a regulatory / policy decision y from a policy space $Y = \{-1, 1\}$. Intuitively, Y may embody a choice between a politically “Conservative” policy ($y = 1$) and a “Liberal” one ($y = -1$). For example, actors in the hierarchy may be considering whether to preserve a deregulatory *status quo ante* (such as not requiring passive safety restraints in passenger cars) or to adopt a regulatory intervention (requiring them).¹⁸

Although actors are assumed to have *a priori* policy preferences (see below), we also suppose that they care about the *fit* between the ultimate policy choice and an objective state of the world – what we will call *facts*. In the example above, facts might embody information about how effective passive restraints are relative to their costs. We presume that a random variable $X \in \mathbb{R}$ represents the true facts, and that X is commonly known *ex ante* to be normally distributed with mean μ and precision τ .¹⁹ (Our framework also admits the limiting degenerate case when $\tau \rightarrow 0$, so that priors are essentially uninformative).

3.1.1 Judicial and Agency Actions and Preferences

The true realization of facts, x , is important to all decision makers because it affects their ultimate policy preferences. In particular, we assume that each regulatory / judicial actor i faces quadratic payoffs over policy outcomes of the form $-(x + \theta_i - y)^2$, where x and y are as described above, and $\theta_i \in \mathbb{R}$ denotes actor i ’s political predispositions – or *ideology* – which we assume to be drawn independently from common distribution $H(\theta)$. We place little structure (at this stage) on the nature of this distribution, though a common assumption in the literature is that it consists of two mass points, corresponding to Democrats ($\theta_i = \theta_D$) and Republicans ($\theta_i = \theta_R > \theta_D$). In any event, our structure implies that each actor possesses an ideal point in policy space, $y_{\theta_i}^* = x + \theta_i$, with utility

¹⁸Our framework easily extends to allow for more than two outcomes. We address this extension in Section 5.

¹⁹Because normal distributions make our analysis more tractable, we utilize them throughout the analysis below. As will become clear, however, our general arguments do not turn crucially on this distributional form.

falling in the squared distance from that point. Note that while actors' political inclinations (θ_i) remain fixed, the locations of their ideal points also depend on realized facts (x). Our framework therefore allows for actors who may lean left or right on *a priori* grounds, but need not be dogmatic ideologues.²⁰ In principle, facts could be strong enough to overcome ideology, inducing a (say) liberal judge/agency to favor a conservative policy (or vice versa). Such persuadability – at least for the median voter – seems central to the deliberative process.

Figure 1: Ideal point as a function of facts (x) & ideology (θ)

Figure 1 illustrates the ideal point mapping for the specific case where $x = \frac{1}{2}$, comparing the ideal point of two decision makers: a Democrat (with $\theta_i = \theta_D \equiv -1$); and a Republican (with $\theta_i = \theta_R \equiv 1$). In the figure, when the Republican judge (who leans toward conservative policies) observes a relatively conservative set of facts ($x = \frac{1}{2}$), her ideal point is relatively conservative, at $y_R^* = 1.5$. If constrained to choose policy $y \in \{-1, 1\}$, then, she will clearly prefer $y = 1$. The Democrat, in contrast, leans liberal; observing the same facts pushes her mildly right, but only enough to move her ideal point to $y_D^* = -0.5$. Thus, the Democrat judge would continue to favor $y = -1$, but with more ambivalence about her position than her Republican counterpart. Were x to take on a

²⁰ Our model treats ideology as a primitive. We therefore do not address what might generate heterogeneous ideologies to begin with; nor do we consider whether information aggregation can cause ideologies to converge. See, e.g., Aumann (1976).

more extreme realization ($x > 1$), it would be enough to sway the Democrat to support the conservative outcome. (And symmetrically with the Republican for $x < -1$).

Our model injects a significant complication into the story illustrated by Figure 1. Specifically, decision makers in this model never know with certainty what the “true” facts are. Rather, they must attempt to maximize their expected utility from a chosen policy, given the information that is available (described in greater detail below).

The judicial review process in our game consists of two stages. In the first stage, the lower level actor (Player A) possessing ideology θ_A makes a decision about legal/regulatory policy. In reaching its decision, A is privy to a signal $Z \in \mathbb{R}$, which conveys noisy information about x . We assume Z is normally distributed with unbiased mean x and finite precision γ . (We also assume that this signal is either collected at no incremental cost, or its collection is non-discretionary to A). After observing the signal, Player A announces a regulatory rule, $y = -1$ or $y = 1$.

Upon Player A ’s announced decision, the second stage begins, where with probability $\pi \in (0, 1)$ an appeals court is asked to review A ’s initial ruling.²¹ The appellate court, denoted collectively by J , is composed of an odd number of $(2M - 1)$ judges, where $M \in \{1, 2, 3, \dots\}$.²² Its panelists are chosen at random from the judiciary pool once an appeal occurs. Let $\Theta \equiv \{\theta_1, \dots, \theta_{2M-1}\}$ denote the set of judicial ideologies on the panel. Without loss of generality, one can re-index the panelists in terms of their ideological order statistics, $\{\theta_{(1)}, \dots, \theta_{(2M-1)}\}$, so that $\theta_{(1)}$ corresponds to the ideology of the most liberal judge on the panel, and $\theta_{(2M-1)}$ corresponds to the ideology of the most conservative judge. We will be particularly interested in the 3-tuple $\hat{\Theta} \equiv \{\theta_{(1)}, \theta_{(M)}, \theta_{(2M-1)}\}$, which includes the most liberal, the most conservative, and median ideologies of the panel (its “Left, Right, and Center” as it were).

Should the appellate panel hear the case, it costlessly observes the realization of Z – the factual signal / record upon which the agency relied. In addition, however, each panelist may individually invest in an auditing technology that generates an additional signal V about the facts. We assume V to be independent of Z , and normally distributed with unbiased mean x and precision ρ . Collecting this signal imposes a cost $c > 0$ on the auditing judge (which reflects effort, docket pressures, opportunity costs, and so forth).²³ For the moment, we assume that each panelist acts independently in deciding whether to audit, and that the signal V is a common value across panelists. Thus, once one judge audits, no new information is provided by additional observations of V . If any judge(s) purchase V , we suppose for now that the additional signal

²¹We assume π to be exogenous. It can be endogenized in a more complex model. E.g., Cameron and Kornhauser (2006).

²²A three-judge panel, therefore, would correspond to $M = 2$; the U.S. Supreme Court would correspond to $M = 5$.

²³All our results carry over easily to the case where the realized value of c is stochastic, drawn from a distribution function $G(c)$ defined on $c \in (0, \infty)$.

is revealed immediately to other members of the panel.²⁴ The panel thereupon makes a decision by majority vote. Should the panel overturn A’s decision, we suppose that A suffers a reputational cost equal to $\varepsilon \geq 0$. To characterize a solution for this game, we require that all players’ policy votes and auditing decisions are consistent with Bayesian perfection.

3.1.2 Motivating J’s Extra Signal

Before proceeding, we pause briefly to motivate our assumption about an additional “signal” available to members of J through auditing. What does it mean, in institutional terms, for an appellate court panelist to spend significant resources to take another draw on the facts? The simplest interpretation would mean a closer examination of the materials in the docket. Although those materials are usually the same ones that the agency / trial judge considered, appellate judges (and their clerks) have distinct backgrounds and abilities, and they act with the benefit of a time lag. Such distinctions may justify viewing intensive appellate review as a form of independent “resampling” from the underlying facts.

Alternatively, a reviewing judge could spend resources finding precedents and doctrinal developments that the agency failed to consider, but which would bear on the ultimate outcome. Similarly, a reviewing judge could spend resources working out how the agency’s decision might yield counterintuitive policy effects, either on the issue directly in front of the agency, or on closely connected issues. Under the right circumstances, this type of research might push other judges to change their votes, and is fully consistent with our model.

Third, the appellate court’s subsequent draw may sometimes represent a review of materials that were simply not considered by Player A. For example, amicus briefs often contain or refer to studies that were not considered by the agency/trial court. Agencies in particular often receive studies and written testimony after the closing date for the submission of evidence. Sometimes these studies were being created, but were not yet complete, at the time the Agency closed the docket. In other circumstances studies are done in response to the Agency’s “concise statement of basis and purpose” published in the Federal Register.²⁵ The reviewing court may have some discretion about whether to consider the new materials and how much attention to give to them. For example, in *Scenic Hudson Preservation Conference v. Federal Power Commission*,²⁶ the Second Circuit Court of Appeals reviewed (and unanimously remanded) the FPC’s decision to grant permission to Consolidated Edison to build a pumped storage hydroelectric power plant on the Hudson River. A plausible reading of the opinion is that the court was persuaded (in part) by the petitioners’ introduction of studies – which were not reviewed by the FPC – about the feasibility

²⁴In Section 5, we address how relaxing many of the above assumptions would affect our results.

²⁵Administrative Procedure Act § 553.

²⁶354 F.2d 608 (2nd Cir. 1965), dismissed, 453 F.2d 463 (2nd Cir. 1971), cert. denied, 407 U.S. 926 (1972).

of more environmentally friendly alternatives. Viewed through the lens of our framework, a decision to take another draw could reflect a decision to consider materials submitted after the agency’s / trial court’s docket closed.

3.2 Appellate Review Stage

The first task for characterizing the equilibrium of this game is to analyze the incentives of the members of a representative judicial panel that is hearing an appeal, assuming that A has already rendered a decision. Ultimately, the members of that panel must decide first whether to collect additional information (audit), and second how to vote. To make predictions about their optimal strategies, we must first compare the expected payoffs of informed and uninformed judges, respectively.

3.2.1 Uninformed Preferences and Decisions

Let us begin with a representative uninformed judge with ideology θ_i , who observes Player A ’s signal, z , but has no additional information. Define such an actor’s preferred outcome here to be y_i^U . It is straightforward to confirm that $y_i^U = 1$ (i.e., panelist i favors the conservative outcome) if and only if:²⁷

$$z \geq z_i^U \equiv -\frac{\theta_i(\tau + \gamma) + \tau\mu}{\gamma} \quad (1)$$

It clear by inspection that z_i^U is strictly decreasing in θ_i , and thus for any two decision makers j and k with $\theta_j < \theta_k$, $z_j^U > z_k^U$. Intuitively, the more liberal the panelist, the “harder sell” she is on the conservative outcome. That is, liberal panelists require a higher public signal z than do relatively conservative players in order to support the conservative policy choice. By the same reasoning, conservative actors are a harder sell on the liberal outcome. Should the judicial panel hear the case, of course, its collective decision will track the median voter’s preferences. Consequently, the uninformed panel’s decision will correspond to the median voter’s preferred outcome, y_M^U , so that the majority favors the conservative over the liberal outcome if and only if $z \geq z_M^U \equiv -\frac{(\theta_{(M)}(\tau + \gamma) + \tau\mu)}{\gamma}$.

After some algebraic manipulation, one can show that a panelist with ideology θ_i sitting on an uninformed panel will realize an expected payoff of:

$$\begin{aligned} \pi_U(\theta_i|z, \theta_{(M)}) &= -E_{x|z} \left\{ ((x + \theta_i) - y_M^U)^2 | z \right\} \\ &= -\left(\frac{1}{\tau + \gamma} + \left(\frac{\tau\mu + \gamma z}{\tau + \gamma} + (\theta_i + 1) \right)^2 \right) + \begin{cases} 0 & \text{if } z \leq z_M^U \\ 4 \left(\theta_i + \frac{\tau\mu + \gamma z}{\tau + \gamma} \right) & \text{else} \end{cases} \end{aligned} \quad (2)$$

²⁷The expression emerges from the observation that $(X|Z)$ is normally distributed with mean $\frac{\tau\mu + \gamma z}{\tau + \gamma}$, and precision $\tau + \gamma$. A number of the other derivations below also depend on manipulated Gaussian distributions. See the appendix for details.

The intuition behind this payoff structure is perhaps best understood through a numerical example. Consider Figure 2 below, for the parametric case where $\mu = 0$, $\tau = 0.5$, and $\gamma = \rho = 1$. The figure envisions a three-judge panel consisting of a liberal ($\theta_{(1)} = -1$) a centrist ($\theta_{(2)} = 0$) and a conservative ($\theta_{(3)} = 1$), and depicts for each judge the expected payoffs associated with both the liberal and the conservative policy choices (the black and gray curves, respectively). In addition, each curve distinguishes between equilibrium payoffs (solid lines) and out-of-equilibrium payoffs (dashed lines). In Figure 2B, corresponding to the centrist panelist, note that the judge's equilibrium payoff corresponds to the outer hull of these curves, reflecting the power of the median voter to dictate outcomes. So long as the panel remains uninformed, its decision will track the median judge's preferences as illustrated in Figure 2B. Note also that a local minimum of the median judge's expected payoff occurs at $z_M^U = 0$, where she is indifferent (or perhaps more accurately, ambivalent) between the conservative and liberal policies. In Figure 2A, the liberal panelist's predispositions imply that it takes a relatively strong factual case ($z > 1.5$) to sway her to favor the conservative policy. Consequently, her equilibrium payoff experiences a downward discontinuity at $z = 0$, corresponding to the fact that at this point the median panelist would swing over to the conservative policy outcome (prematurely, from the liberal judge's perspective). Figure 2C illustrates the opposite case – a conservative panelist whose ideology is $\theta_i = 1$. For this judge the indifference point between outcomes occurs at $z_{-1}^U = -1.5$, reflecting the fact that it takes an analogously strong case ($z < -1.5$) to sway the conservative actor to favor the liberal policy. Similar to the liberal panelist, the conservative judge's payoff also realizes a discontinuity (this one upward) at $z = 0$, reflecting the point where the median voter begins to favor the conservative policy.

Figure 2. Uninformed Expected Payoffs of Mixed Judicial Panelists

As will become evident below, the location of the median judge's indifference point – and the corresponding payoff discontinuities for the non-median judges – relate directly to auditing incentives within the panel.

3.2.2 Informed Preferences and Decisions

Now consider the case where at least one panelist audits, so that the representative judge i with ideology θ_i observes both z and v . Define an informed actor's

preferred policy choice to be y_i^I . It is straightforward to confirm that $y_i^I = 1$ (i.e., panelist i favors the conservative outcome) if and only if:

$$v \geq v_i^I \equiv - \left(\frac{\theta_i \cdot (\rho + \tau + \gamma) + z\gamma + \tau\mu}{\rho} \right)$$

In other words, an informed judge favors the conservative option over the liberal one whenever the auditing signal, v , is sufficiently strong relative to her ideology, her priors about x , and the content of A 's original signal, z . Similar to the uninformed panel, an informed panel will issue a holding coinciding with the informed median judge's preferred outcome, or y_M^I . Therefore, an informed panel opts for the conservative outcome if and only if $v \geq v_M^I \equiv - \left(\frac{\theta_{(M)} \cdot (\rho + \tau + \gamma) + z\gamma + \tau\mu}{\rho} \right)$.

Consequently, a panelist with ideology θ_i sitting on an informed panel realizes an expected payoff of:

$$\begin{aligned} \pi_I(\theta_i|z, \theta_{(M)}) &= -E_{v|z} \left\{ E_{x|z,v} (x + \theta_i - y_M^I)^2 | z, v \right\} \\ &= - \left(\frac{1}{\tau + \gamma} + \left(\frac{z\gamma + \tau\mu}{\tau + \gamma} + (\theta_i + 1) \right)^2 \right) \\ &\quad + 4 \cdot \left(\theta_i + \left(\frac{\tau\mu + \gamma z}{\tau + \gamma} \right) \right) \left(1 - \Phi \left(- \frac{\left(\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma} \right)}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}} \right) \right) \\ &\quad + 4 \cdot \sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}} \cdot \phi \left(- \frac{\left(\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma} \right)}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}} \right), \end{aligned} \quad (3)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ represent the standard normal probability density and cumulative distribution functions, respectively.

3.2.3 The Value of Information

Having characterized the expected payoffs associated with both uninformed panels and informed panels, we are now in a position to consider the difference – denoted as $\Delta(\theta_i|z, \theta_{(M)})$ – between the judge's expected payoff in the informed state and its counterpart payoff in the uninformed state. Implicitly, $\Delta(\theta_i|z, \theta_{(M)})$ corresponds to the equilibrium value that each judge places on

additional information (in the form of signal v). Subtracting (2) from (3) yields:

$$\begin{aligned}
\Delta(\theta_i|z, \theta_{(M)}) &\equiv \pi_I(\theta_i|z, \theta_{(M)}) - \pi_U(\theta_i|z, \theta_{(M)}) \\
&= 4 \cdot \sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}} \cdot \phi\left(-\frac{\left(\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}\right)}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) \\
&\quad + 4\left(\theta_i + \frac{z\gamma + \tau\mu}{\tau + \gamma}\right) \cdot \begin{cases} \left(1 - \Phi\left(-\frac{\left(\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}\right)}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right)\right) & \text{if } z \leq z_M^U \\ -\Phi\left(-\frac{\left(\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}\right)}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) & \text{if } z > z_M^U \end{cases}
\end{aligned} \tag{4}$$

Despite its apparent complexity, this expression generates a number of core intuitions from this paper. Note that for each judge i , the value of information hinges on both the judge's own ideology (θ_i) and that of the median panelist ($\theta_{(M)}$). This makes sense, since the judge's own policy commitments should factor into whether she finds more information helpful, but so should the pragmatic assessment of how additional information may affect the ultimate outcome – by swaying the median judge. Harvesting another signal, therefore, may not only help to refine the auditing judge's assessment of her personally preferred policy, but it may also help her persuade a recalcitrant median judge to join her. Alternatively, a judge may place little or no value on auditing when the median judge already is leaning in her direction, since added information may backfire, persuading the median judge to withdraw her allegiance. Consequently, strategic, non-median panelists are relatively reluctant to audit when doing so materially imperils the support of the median voter. These intuitions are stated formally in the following two lemmas:

Lemma 1: *For the median judge with ideology $\theta_{(M)}$, auditing is maximally valuable at her uninformed indifference point, $z = z_M^U$, and falls symmetrically in both directions as z diverges from z_M^U .*

Lemma 2: *If judge i is more conservative than the median judge ($\theta_i > \theta_{(M)}$):*

- Judge i values information more than the median judge when $z \leq z_M^U$ and less than the median judge when $z > z_M^U$.
- The extent to which the more conservative judge's valuation exceeds / falls short of the median judge's increases strictly in θ_i .

If judge i is more liberal than the median judge ($\theta_i < \theta_{(M)}$):

- Judge i values information more than the median judge when $z \geq z_M^U$ and less than the median judge when $z < z_M^U$.
- The extent to which the more liberal judge's valuation exceeds / falls short of the median judge's decreases strictly in θ_i .

The intuitions behind Lemmas 1 and 2 are perhaps best understood through a numerical example. Figure 3, below, returns to the same calibration as Figure 2, involving a three judge panel composed of a liberal judge, a centrist median judge, and a conservative judge, so that $\Theta = \{-1, 0, 1\}$, and $\mu = 0$, $\tau = 0.5$, $\gamma = \rho = 1$. Each diagram represents the value that the liberal, moderate, and conservative judge attaches – in equilibrium – to the additional signal v , as a function of the agency’s initial signal z . The median judge (Figure 2B) always places positive value on the extra signal, since she dictates the final outcome, and such information can only help her refine this choice. In fact, an additional signal is *most* valuable when $z = 0$ – the point where the median panelist is maximally ambivalent between the liberal and conservative policy options. As z moves away from this point of indifference, her preferred policy choice becomes more clear cut, and in turn the value she places on additional information decays (symmetrically, as noted in Lemma 1).

Figure 3. Expected Information Value of Auditing for Mixed Judicial Panelists

In contrast, the liberal and conservative judges (Figures 3A and 3C, respectively) attach more complicated equilibrium valuations to additional information (as described in Lemma 2). The liberal judge, for example, values additional information only when the agency’s signal $z \geq z_M^U = 0$. Moreover, in this region, the liberal judge places a much higher value on learning the new signal than either of the other panelists. When $z < 0$, in contrast, the liberal judge may expect to suffer *disutility* from additional information, and in any event places a much lower value on information than the other panelists. The intuition for this result is as follows: when $z \geq 0$, the liberal judge knows that absent more information, the median panelist will support the conservative policy outcome. If she is able to convince the median judge to switch sides, the liberal judge can expect to receive a discontinuous upward shock to her payoff. But she cannot win over the median judge without some deliberative ammunition; by auditing, she may discover information that will bring the median voter on board, and in the process enhance her payoff. In contrast, when $z < 0$, the median panelist is already leaning towards the liberal policy; additional information, while nice in the abstract, runs an appreciable risk of pushing the median panelist back across the political aisle. In the example pictured in Figure 3, this latter threat

is so significant that it swamps other plausible values from auditing when $z < 0$ for the liberal panelist. Exactly the opposite logic follows for the conservative judge: she places significant value on auditing when $z \leq z_M^U = 0$, so that the median judge initially favors towards the liberal outcome. In contrast, the conservative judge places no value (indeed negative value) on more information when $z > 0$.

All told, in this example *either* the liberal *or* the conservative judge (but generally not both) has the greatest incentive on the panel to collect additional information. As it turns out, this logic carries over more generally to panels of arbitrary size and ideology, an insight stated formally in the following lemmas:

Lemma 3: *When $z < z_M^U$ the most conservative judge (with ideology $\theta_{(2M-1)}$) has the maximal incentive of all panelists to audit. When $z > z_M^U$, the most liberal judge (with ideology $\theta_{(1)}$) has the maximal incentive to audit. If $z = z_M^U$, the most conservative (most liberal) panelist has the greatest incentive to audit when $(\theta_{(2M-1)} - \theta_{(M)})$ is larger (smaller) than $(\theta_{(M)} - \theta_{(1)})$.*

Lemma 4: *If $c \leq c(\hat{\Theta}, z)$, at least one panelist has an incentive to audit (and thus learn v), where*

$$c(\hat{\Theta}, z) = 4\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}} \cdot \phi\left(-\frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) \quad (5)$$

$$+ \begin{cases} 4\left(\theta_{(2M-1)} + \frac{\tau\mu + \gamma z}{\tau + \gamma}\right) \cdot \left(1 - \Phi\left(-\frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right)\right) & \text{if } z \leq z_M^U \\ -4\left(\theta_{(1)} + \frac{\tau\mu + \gamma z}{\tau + \gamma}\right) \cdot \Phi\left(-\frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) & \text{if } z > z_M^U \end{cases}$$

This condition implicitly defines an auditing interval $\left[\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})\right]$ around z_M^U .

Note from Lemma 4 that the auditing interval is completely characterized by the ideologies of the median judge and the two ideologically extreme judges. No other judge's ideology enters into the expression from Lemma 4. In general, as the ideological wings of the panel grow more extreme, so too does the auditing interval (and with it grows the prospect of reversal).

A number of corollaries immediately follow from inspection and/or differentiation of the expression in Lemma 4:

Corollary 4.1: *The auditing interval $\left[\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})\right]$ is strictly increasing in the precision of the auditing technology (ρ), and strictly decreasing in the precision of A 's signal (γ) and in the realized cost of auditing (c).*

Corollary 4.2: *The auditing interval is strictly increasing in $\theta_{(2M-1)}$ and strictly decreasing in $\theta_{(1)}$.*

Corollary 4.3: *The auditing interval is potentially asymmetric.*

Corollary 4.4: *The auditing interval is invariant to all median- and extrema-preserving transformations of Θ .*

Corollary 4.1 is intuitive. Corollary 4.2 embodies the idea that all else constant, greater panel diversity (as measured by the ideologies of the most ideologically extreme panelists) is more likely to induce informed scrutiny of the agency’s decision. This effect emanates directly from the fact that preference differences between the median voter and the extreme wings of the panel are what drive the latter to audit when the median would not. Amplifying those ideological differences (i.e., increasing $\theta_{(2M-1)}$ or decreasing $\theta_{(1)}$) enhances this effect. Corollary 4.3, however, suggests that diversity need not inculcate symmetric scrutiny. In particular, as the conservative (liberal) wing of the party becomes more distinct from the median, the panel is increasingly likely to reject liberal (conservative) policies that the median voter would have favored if uninformed; but it is no more or less likely to reject conservative (liberal) policies that the uninformed median voter would have favored. Finally, Corollary 4.4 formally restates the result that the auditing interval turns solely on ideologies of the Left, Right, and Center judges. Consequently, holding those ideologies constant, our model predicts identical auditing ranges (and reversal rates) for a 3-judge, 5-judge, 9-judge, or even a 99-judge panel.²⁸

3.3 Player A’s Optimal Strategy

Having characterized the continuation payoffs of the judicial panel J conditional on appeal, we now move backwards in sequence to analyze the strategy of Player A (the agency / trial court), who anticipates the equilibrium strategy described above. Recall that A is motivated both by a desire to implement her preferred outcome and to avoid being overturned. Moreover, recall that with probability $(1 - \pi)$, Player A ’s decision will never be appealed, in which case the best she could do is to implement her sincere policy choice given z . On the other hand, if A ’s decision is appealed (with probability π), her payoff becomes more complicated. On the one hand, A suffers a cost ε should her decision be overturned. But on the other hand, if the reviewing court also augments A ’s information through judicial review, it will issue a more informed policy choice, which will also affect – and possibly increase – A ’s welfare. Combining these factors, A ’s expected payoff given z is:

²⁸We should note that generalizations of our model could weaken this invariance result. For example, if panelists faced differential costs in auditing, a moderate judge with particularly low auditing costs may place a higher *net* benefit on auditing than an extreme judge who faces a high cost of auditing. Similarly, if a relatively moderate judge can collect a significantly more precise signal than an extreme judge, the former may determine the extreme end of the auditing range.

$$\begin{aligned}
& (1 - \pi) \cdot E_x \left(- (x + \theta_A - y_A)^2 \mid z \right) \\
& + \pi \cdot \left(\begin{aligned} & (1 - q_{\hat{\Theta}, z}) \cdot E_{\hat{\Theta}, x} \left(- (x + \theta_A - y_M^U)^2 - \begin{cases} 0 & \text{if } y_A = y_M^U \\ \varepsilon & \text{else} \end{cases} \mid z \right) \\ & q_{\hat{\Theta}, z} \cdot E_{\hat{\Theta}, x, v} \left(- (x + \theta_A - y_M^I)^2 - \begin{cases} 0 & \text{if } y_A = y_M^I \\ \varepsilon & \text{else} \end{cases} \mid z, v \right) \end{aligned} \right)
\end{aligned} \tag{6}$$

where $q_{\hat{\Theta}, z}$ is an indicator function taking on a value of 1 if, for an ideology configuration of $\hat{\Theta}$, the agency's signal z lies within the panel's auditing range.²⁹ Remember that because panelists are not selected until *after* Player A has reached a decision, she does not know for sure whether the facts of the case before her will fall within J 's auditing interval, and she must therefore form expectations over the probability density of the ordered three-tuple $\hat{\Theta} = \{\theta_{(1)}, \theta_{(M)}, \theta_{(2M-1)}\}$, which we denote as $f(\hat{\Theta})$.³⁰

Inspection of (6) allows some simplification of its analysis. First, note that A 's decision, y_A only enters into this expression in two ways: (1) It directly affects A 's payoff in the event that no appeal is heard, and (2) it affects the potential costs that A may suffer if she is overturned by J . Numerous other terms of this expression, including $q_{\hat{\Theta}, z}$, $(x + \theta_A - y_M^U)$ and $(x + \theta_A - y_M^I)$, are invariant to A 's ultimate choice and can effectively be held constant. These observations, in turn, yield Lemma 5:

Lemma 5: *Given the equilibrium behavior of a panel with configuration $\hat{\Theta}$, A will favor the conservative outcome if and only if:*

$$4(1 - \pi) \left(\frac{\tau\mu + \gamma z}{\tau + \gamma} + \theta_A \right) \geq \pi \varepsilon \cdot \left(\begin{aligned} & E(\Pr\{z < \underline{z}(\hat{\Theta})\} \mid z) - E(\Pr\{z > \bar{z}(\hat{\Theta})\} \mid z) \\ & + E\Pr\{z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})] \cap v < v_M^I \mid z\} \\ & - E\Pr\{z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})] \cap v \geq v_M^I \mid z\} \end{aligned} \right) \tag{7}$$

To best understand the expression in Lemma 5, it helps to think about the limiting cases for the probability of appeal (π) and the cost to being reversed (ε). Consider first how the agency would decide a case if it was indifferent to reversal, so that $\varepsilon \approx 0$. This assumption may be plausible in many situations. One could imagine, for example, that agency appointees are not sufficiently long lived to worry about reversals months or years later; or that public attention tends to wane over time so that later reversals have little political salience; or that agencies derive considerable utility from expressing their policy stance (independent of the policy's eventual implementation). In such environments, the right hand side of (7) effectively disappears, and A 's choice boils down to selecting the conservative outcome if and only if the left hand side of this expression

²⁹In terms of Lemma 4, $q_{\hat{\Theta}, z} = 1$ iff $z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})]$.

³⁰The functional form of $f(\hat{\Theta})$ is provided in the Appendix.

is nonnegative, which occurs when $z \geq z_A^U$. Not surprisingly, this limiting case is identical to having the agency implement its sincere policy preference.

Similarly, suppose the cost of being overturned were nontrivial, but that the probability of appeal were negligible, so that $\pi \approx 0$. Here once again, the agency would focus on the left hand side of the expression above, and it would generate a sincere policy decision.

In contrast, suppose that both the costs of reversal were nontrivial and that the probability of appeal were close to unity. Here, the agency knows that its initial policy choice is almost certainly going to be revisited, and thus the left hand side of the above expression carries little significance. Player A will therefore focus on the right hand side of (7), favoring a conservative policy if:

$$\begin{aligned} & \underbrace{\Pr\{z < \underline{z}(\hat{\Theta})|z\} + E \Pr\{z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})] \cap v < v_M^I | z\}}_{E(\Pr\{J \text{ reverses a conservative decision } | z\})} \quad (8) \\ \leq & \underbrace{\Pr\{z \geq \bar{z}(\hat{\Theta})|z\} + E \Pr\{z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})] \cap v \geq v_M^I | z\}}_{E(\Pr\{J \text{ reverses a liberal decision } | z\})} \end{aligned}$$

Although this expression looks complicated, its interpretation is simple – Player A 's strategy devolves into minimizing the expected probability of reversal. Effectively, the agency attempts to mimic the most likely decision of the future reviewing panel.

Finally, consider intermediate cases, where ε and π are large enough to matter but not so large as to dominate. Here, Player A cares about both its sincere policy commitments and its aversion to reversal. Consequently, the A 's behavior will turn on the relative stakes from each concern. For example, if it observes a value of z very close its indifference point, z_A^U , then Player A will have only weak preferences over available outcomes. Accordingly, its desire to avoid reversal will predominate. On the other hand, suppose the distribution of judicial ideologies – from which appellate judicial panels are drawn – is highly dispersed, making it prohibitively difficult to make reliable predictions about the panel's ultimate composition. Here, Player A may rationally throw up its hands at the prospect of anticipating future decisions, concentrating instead on issuing a sincere policy decision.

3.4 Equilibrium

Having characterized the continuation payoffs of both A and all panelists in J , we are almost in a position to state the equilibrium for the game. Before doing so, however, it is necessary to attend to a technical issue relating to equilibrium selection for the panelists on J . As should be clear from the above discussion, there may frequently be cases where more than one judge on a panel is willing to audit. Because auditing provides a common informational good to all, auditing by more than one panelist is not part of a pure strategy equilibrium, and any mixed strategy equilibria that support such outcomes are easily dominated by

numerous coordinated pure strategy equilibria. Thus, it is sensible to assume that the panelists will find some mechanism for coordinating their investments. One such mechanism, which we presume hereafter, is as follows:

Assumption A: (1) *If multiple judges on the same panel value additional information enough to justify auditing, then the judge who places the maximal value on the additional signal is presumed to bear the cost of auditing.*
(2) *If two or more judges on the same panel share the maximal value of an additional signal, they randomize as to who audits.*

Although Assumption A seems reasonable (at least to us), there are many alternatives that would generate outcome-equivalent equilibria.³¹ Applying this selection assumption to the Lemmas above, the following result immediately emerges:

Proposition 1: *If Assumption A holds, the following is the unique equilibrium of the auditing game:*

- *The agency issues a conservative opinion iff the condition in (7) is satisfied;*
- *If an appeal occurs, and if $z \in [\underline{z}(\hat{\Theta}), z_M^U]$, panelist $\theta_{(2M-1)}$ audits (revealing v), and the panel issues a conservative decision iff $v \geq v_M^I$;*
- *If an appeal occurs, and if $z \in (z_M^U, \bar{z}(\hat{\Theta})]$, panelist $\theta_{(1)}$ audits (revealing v), and the panel issues a conservative decision iff $v \geq v_M^I$;*
- *If an appeal occurs, and if $z = z_M^U$, the extreme panelist $\theta_{(1)}$ or $\theta_{(2M-1)}$ that is furthest from $\theta_{(M)}$ audits (revealing v). If both are equidistant from $\theta_{(M)}$, they randomize as to who audits. The panel issues a conservative decision iff $v \geq v_M^I$.*
- *If an appeal occurs, and if $z \notin [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})]$, the panel issues a conservative policy (overturning Player A if necessary) iff $z \geq z_M^U$.*

As with Lemma 4, in Proposition 1 the auditing decisions and policy choice of the panel are fully characterized by the ideologies of the Left, Right and Center panelists. No other judge's ideology enters into the expression from Proposition 1 (at least with this characterization of the model). Notwithstanding the dominance of the median voter model in positive political theory, then, the results above suggest ways in which judicial panels (and other deliberative bodies) respond to their extreme wings rather than the middle. As such, it joins a growing literature in documenting how non-median members can affect outcomes, by lobbying, influencing, shaming, or (in our case) altering the course of endogenous information production.

³¹For example, an alternative assumption (that is outcome equivalent) posits that judge i audits a case with initial signal z if (1) she places a positive net value on auditing, and (2) the next judge closer to the median judge (if she exists) does not place a positive net value on auditing.

4 Examples, Simulations, and Empirical Implications

This section teases out some of the immediate implications of our model for the empirical literature on judicial panel effects. As noted in Section 2, this literature provides significant support for the moderating effect of ideological diversity, documenting a tendency of minority and majority judges on mixed panels to move towards each other when voting. Below we demonstrate how this (and other) predictions can play out in our model, first in the form of a numerical example, and then in a simulation calibrated against real-world data. Finally, we develop some preliminary thoughts about how our model might be tested against alternatives in the literature.

4.1 Numerical Example

Consider a numerical example of our model involving a three-judge panel. Suppose that the agency is a Democrat ($\theta_A = -1$), that the cost of reversal (ε) is negligible, and – as in the Figures above – that $\tau = 0.5$, $\mu = 0$, and $\gamma = \rho = 1$. At these parameter values, it is straightforward to confirm that the equilibrium probabilities of a liberal and conservative decision by the agency (respectively) are 80.7% and 19.3%.

4.1.1 Homogenous Democrat Panel

Consider first a judicial panel composed entirely of Democrat ideologies, so that $\theta_{(1)} = \theta_{(M)} = \theta_{(3)} = -1$, which we denote as a DDD panel. The solid line in Figure 4 illustrates – as a function of prior signal z – the expected value to each panelist of collecting an additional signal v . Notice that the value of information is symmetric around a maximum at $z = 1.5$, which is the point where the Democrat actors are most ambivalent between the policy outcomes. This makes intuitive sense, since the point of maximal ambivalence is where additional information is most useful. In contrast, when $z < -0.5$ or $z > 3.5$, the agency’s signal (z) is so strong that the value of additional information (through v) is negligible.

Figure 4: Auditing Range of DDD Panel

Suppose that the cost of auditing for each judge is $c = \frac{1}{10}$ (represented by the horizontal dashed line). As Figure 4 illustrates, each judge places nonnegative *net* value on auditing only where the agency's signal z falls in the auditing interval $[0.5155, 2.4845]$ (approximately). Inside this interval, one judge is randomly selected to audit, and the panel bases its decision on (z, v) . Outside it, no one audits and the panel bases its decision solely on z . Note that this interval remains symmetric around $z = 1.5$. The resulting equilibrium has the following characteristics:

- The Democrat agency A issues the conservative policy whenever $z \geq 1.5$. Otherwise it issues the liberal policy.
- The DDD panel audits A 's decision whenever $z \in [0.5155, 2.4845]$. Otherwise it rubber stamps A 's decision.
- If the DDD panel audits, it favors the conservative outcome (overturning A if necessary) whenever $v + z \geq \frac{5}{2}$. Otherwise it favors the liberal outcome (overturning A if necessary).
- Viewed ex ante, the DDD court (unanimously) overturns liberal policy positions by the D-agency at a rate of approximately 6.2%. It (unanimously) overturns a conservative policy decision by the agency at a rate of 18.1%. The unconditional rate of reversal of the agency by the DDD panel in this case is 8.5%.

4.1.2 Heterogeneous DDR Panel

Now consider what happens if one replaces a Democrat panelist with a Republican — a DDR panel. Under conventional median voter logic, the injection of a single R panelist should not affect outcomes, since she is not a pivotal voter, and thus the panel's decision rule (i.e., how they translate either z or (z, v) into

policy space y) cannot change from the DDD case, *at least holding information constant*. However, available information may change with the addition of an R panelist, who faces different incentives to become informed of the additional signal (v). In particular, the lone R may wish to audit cases that the majority would not – so long as his inquiry might plausibly sway their opinion. As predicted by Proposition 1, the R judge will tend strategically to audit cases lying solely to the "left" (in z space) of the D majority's indifference points. These are the very cases where the Democrat majority is potentially persuadable, yet may have insufficient incentives to audit acting individually.

The dark solid line in Figure 5 below depicts the maximal valuation that any of the panelists places on auditing (as a function of z). Note that when $z > 1.5$, the diagram is identical to Figure 4. In this region, only the two Democrat judges place a positive value on auditing. The Republican panelist actively resists auditing within this range, since the Democrats are already leaning in his direction, and more information may induce them to reconsider. In contrast, when $z < 1.5$, the Republican is strongly motivated to audit, as reflected by the upward shift of the valuation curve (relative to Figure 4) over that interval.

Figure 5: Auditing Range of DDR Panel

Because of the added motivation of the R whenever $z < 1.5$, the auditing interval for the DDR is (approximately) $z \in [-0.2615, 2.4845]$, representing a leftward expansion from the DDD case ($z \in [0.5155, 2.4845]$). The full equilibrium for the DDR panel is characterized as follows:

- The Democrat agency A issues the conservative policy whenever $z \geq 1.5$. Otherwise it issues the liberal policy.
- The DDR panel audits A 's decision whenever $z \in [-0.2615, 2.4845]$, which expands the DDD's auditing interval asymmetrically to the left. Otherwise it rubber stamps A 's decision.
- If the DDR panel audits, it favors the conservative outcome (overturning

A if necessary) whenever $v+z \geq \frac{5}{2}$. Otherwise it favors the liberal outcome (overturning A if necessary).

- Viewed ex ante, the DDR panel (unanimously) overturns a liberal holding by the D-agency at a rate of approximately 7.2% (which exceeds the corresponding 6.2% reversal rate of the DDD panel). The DDR panel overturns a conservative policy decision by the agency (sometimes unanimously and sometimes on a party line vote) at a rate of 18.1% (which is identical to the DDD panel). The unconditional rate of reversal of the agency by the DDR panel in this case is 9.3% (exceeding the 8.5% unconditional rate for the DDD panel).

The example above demonstrates many of the core intuitions from Proposition 1 and Lemmas 1-5. Expected reversal rates increase when one adds *even a single, non-pivotal minority member*, with the effect being driven solely by an enhanced expected frequency with which a *unanimous* panel reverses liberal agency pronouncements. To an outsider, this might look like the inclusion of the R on the panel has made the Ds more collegial, or the R has threatened to blow the whistle on the Ds. But the effect is distinct: simple self-interest in a noncooperative setting can drive an outcome where more information induces greater apparent moderation. In other words, the pivotal D voter isn't becoming "nicer"; rather, she is using more information, which the R panelist has (strategically) provided her.

4.2 Calibrated Simulations

In order to illustrate the plausibility of our model in real-world settings, we calibrated it to the data set developed by Miles and Sunstein (2008). Of particular interest is their Table 2 (Miles and Sunstein 2008, p. 786), which reports rates at which *individual* judges vote to validate an agency's decision, contingent on the other judges on the panel. The first column of the table below lists each possible panel composition; the second lists the type of judge within the panel; and the third lists Miles and Sunstein's (2008) reported empirical validation rate.³² The simulations reported in the fourth column estimate a single agency political ideology (which is slightly liberal), even though the actual data likely come from several different agency types (in terms of ideology).³³ The results, in our opinion, seem quite good. With a not unreasonable set of parameters we can come very close to the affirmance rates observed by Miles and Sunstein; the simulation averaged an error of 1.84% across all six conditions.

³²The table reflects simulated affirmance rates using parametric values that best fit the actual outcomes under a least squares criterion. Other empirical fit criteria generate similar results.

³³Specifically, we estimate a pooled agency ideology of $\theta_A = -0.157$ and an auditing cost of $c = 0.7458$. We continue to assume in this simulation (as above) that the costs of reversal and/or the probability of appeal are modest, so that the agency issues sincere opinions and the distribution of panel types need not be factored into our analysis. This distributional information was not available in the Miles/Sunstein data. Including it would increase our degrees of freedom, so it would only cause our simulated results to improve.

The parameters used for this simulation reflect a true state of the world that is almost neutral ($\mu = -.0227$) in expectation, but has low precision ($\tau = .3677$). In addition, the agency's signal about the true state of the world appears noisy ($\gamma = .0837$) when compared to the reviewing court's signal ($\rho = 1.3397$). We do not claim that our exercise proves that the parameter values represent the "true" characteristics of the courts and agencies. Rather, our exercise shows that *if* the characteristics embedded within the parameter values used in the simulation happened to be true, then we would observe affirmance rates very similar to those observed by Miles and Sunstein.

Panel Composition	Panelist Type	M&S (2000) Validation Rate	Simulated Val. Rate	Abs. Difference Sim. v. Empirical
(D, D, D)	D	0.746	0.7539	0.0079
(D, D, R)	D	0.697	0.6822	0.0148
(D, D, R)	R	0.667	0.6295	0.0375
(D, R, R)	D	0.678	0.6933	0.0153
(D, R, R)	R	0.604	0.6109	0.0069
(R, R, R)	R	0.551	0.5792	0.0282

Table 1: Simulated Affirmance Rates

Having satisfied ourselves that we could replicate their central summary results, we obtained the actual decision-level data from Miles and Sunstein (2008).³⁴ Again, the object was to see how well we could simulate their measured outcomes, but here we hoped to pull apart the Miles and Sunstein data into more granular categories, and then produce a simulation that came close. In particular, we went through the data and accounted for whether the *administrative agency* was under a Democrat or Republican administration, and then controlled for panel composition. In addition, for this exercise we simulated the probability of a panel (i.e. the median voter) voting to overturn the agency. This exercise produced the following chart:³⁵

³⁴We express sincere thanks to Tom Miles and Cass Sunstein for sharing their data with us.

³⁵For this simulation we minimized the sum of the squares of the values in the last column, and the minimizing parameter values are $\mu = -0.7841$, $\tau = 0.05$, $\gamma = 0.0029$, $\sigma = 0.6395$, and $c = 4.786$. We also fit the model using other maximands.

In addition, we experimented with different lag definitions of when an agency becomes associated with the president in charge, given that upon a change in administration, the incumbent agency may have to continue to defend its actions under the previous administration. Such adjustments do not appear to affect these results significantly.

Agency Type	Panel Composition	M&S (2000) Reversal Rate	Simulated Reversal Rate	Absolute Difference
<i>D</i>	(<i>D, D, D</i>)	0	.1325	.1325
<i>D</i>	(<i>D, D, R</i>)	.4375	.2642	.1733
<i>D</i>	(<i>D, R, R</i>)	.3913	.4573	.0660
<i>D</i>	(<i>R, R, R</i>)	.5	.4969	.0031
<i>R</i>	(<i>D, D, D</i>)	.4	.4445	.0445
<i>R</i>	(<i>D, D, R</i>)	.3928	.3636	.0292
<i>R</i>	(<i>D, R, R</i>)	.4528	.3622	.0906
<i>R</i>	(<i>R, R, R</i>)	.25	.3121	.0621

Table 2. Simulations Using Agency Data

Note that this simulation – while still relatively solid – does not "fit" as well as the prior one. There are likely multiple reasons for this. First, there are two places where the Miles and Sunstein data are non-monotonic. With a Democrat agency, theory predicts that as we move from a DDR to a DRR judicial panel, the reversal rate should not fall. Yet, in the Miles and Sunstein data, the reversal rate from this change in judges falls from .4375 to .3913. Similarly, with a Republican agency, our theory predicts that when we move from a DDR to a DRR panel, the reversal rate should not rise. Again, the Miles and Sunstein granular data contradict this prediction. This effect could be mere chance (e.g., the number of cases in some of the cells is fairly small); or relatedly, it might pertain to unobserved characteristics (issue level or judge level) that neither we nor Miles and Sunstein (2008) account for. In any event, when we fit the parameters in our simulation, our model does not allow the reversal rate to fall and rise in such non-monotonic ways in the way that it does in the data. Second, we are trying to fit the data on more dimensions in this simulation than in the prior one, which will tend to produce a looser fit. Third, in these simulations we assign ideology specific scores to the Democrat agency ($\theta = -1$) and the Republican one ($\theta = 1$). In Table 1, in contrast, we instead estimated an agency ideology that we applied across all cases (effectively giving us more degrees of freedom to calibrate). Nevertheless, even with these restrictions our model does reasonably well. The only large divergences take place in the first two rows of the chart, and perhaps in the R-DRR case. The other cases fit quite well.

At least as an initial matter, we find these simulations suggestive. They illustrate that there are parameter values for our model that make it perform, more or less, like existing empirical studies of judges. If our model had failed such a test – if we could not find parameters that made the model resemble observed empirical patterns – then we would regard the model with some skepticism. However, since it passed this initial test at (least from an eyeballing perspective), it should be a serious candidate for testing in future work. While we do not concentrate on it here, another artifact of our model may be consistent with other empirical stylized facts in the panel effects literature. Although

the sole R-panelist in the DDR panel is uncovering information strategically, for the purposes of swaying his D counterparts, it is possible that his additional digging will generate a signal that has the opposite effect: That is, it convinces the R-panelist that the liberal policy outcome is optimal even from his perspective. This effect is a small one, but under some circumstances the additional digging undertaken by the minority panel member can also cause him to switch allegiances.

4.3 Empirical Implications

Although the calibration exercise above illustrates the plausibility of our deliberative theory of panel effects, it is not a test of our theory *per se*. To test it, we would need to isolate situations where our model theory delivers predictions that are different from alternative theories (including certain versions of whistleblowing, social collegiality, or attitudinal drift). Constructing such tests must be done with care (Epstein, et al, 2005; Sisk and Heise 2005; Fischman 2009). Although we leave for future work the task of designing a suitable set of cases for just such a comparison, we offer some possibilities below.³⁶

There are a number of potential approaches for testing of our model empirically. For example, Landes and Posner (2009) find strong evidence of mixed-ideology moderation on three-judge panels, but fail to find it on the Supreme Court. Their failure to detect a moderation trend at the US Supreme Court level may be due to any number of factors. However, our model suggests one possibility – that they measured Court ideology through central tendency measures (e.g., percent Republican-appointed) rather than variation at the extremes of the court’s ideological spectrum. Our analysis suggests that variation at the extremes (e.g., changes in the left-most or right-most wings of the court) are more likely to predict changes in auditing intensity and resulting panel effects.³⁷

We may also gain empirical traction from the fact that our model produces panel effects in environments that are both *information poor* and *politically charged*. That is, limited information affords judges with the opportunity to investigate more, and political differences provide them (or at least some of them) with motivation to do so. Our framework therefore suggests that we are most likely to observe panel effects in domains where both characteristics are present (such as in environmental law, securities regulation, or antitrust), and not in fields that are more purely political (such as abortion or gun control policy) or technocratic (such as weights and measures policy).

Our account may also shed light on the role of merit in the Supreme Court confirmation process. Epstein and Segal (2005) measure merit by coding newspaper editorial evaluations of a nominee. They report that, other things being equal, merit is positively correlated with senatorial votes for the candidate. Per-

³⁶There is a sense in which our contribution already *is* empirically driven – for our starting motivation (see Section 2) is to explain stylized empirical facts that a sizeable literature has already identified.

³⁷Appeals court panels, in contrast, consist of only three judges, and such aggregate measures do a better job of capturing ideological variation at *both* the median and the extremes.

haps surprisingly, this effect is strong enough to overcome all or most sources of political opposition. Thus, Scalia was confirmed unanimously, and Ginsburg was confirmed with only three dissenting votes. Our model provides at least a partial explanation if “merit” can be interpreted as judicial accuracy. In our model, all judges are assumed to be equally precise in harvesting additional facts; ideologically isolated panelists have an incentive to work hard to provide better information to the panel. To be sure, the minority judge’s efforts work in her favor; but perhaps less obvious is the fact that *majority judges may also be better off* by the inclusion of the minority judge, due to the public informational good she provides. This latter effect is accentuated further when the minority judge’s sampling precision increases. In this sense, meritorious ideological outliers can be good for everyone—a possibility that even a callous politician can love (or at least learn to live with).

Finally, our framework may provide an alternative approach for estimating ideological “scores” for judges and other legal actors (e.g., Martin and Quinn 2002; Epstein et al. 2007b). In much of the existing ideological scoring literature, identification is achieved through an attitudinal model of voting that assumes complete information and excludes deliberation. One can estimate ideological scores under our framework too, but identification is based on a deliberative model of voting with incomplete information and endogenous search. Once estimated, the predictive power of these alternative scores could be compared to their attitudinal analogs (e.g., Martin-Quinn scores) as a means for testing the deliberative against the attitudinal model.³⁸

5 Extensions

The analytical framework presented above also lends itself to a number of theoretical generalizations and extensions. Although we do not analyze all of them here, we briefly address some of the more promising ones, noting their likely effects on our model’s predictions.

A first extension of the model involves altering the informational environment at the *review* stage. For instance, one could imagine a structure (following on Spitzer and Talley (2000) and Cameron et al. (2000)), where the appellate panel cannot directly observe the factual input (z) that undergirds A ’s policy choice, but instead makes equilibrium inferences from A ’s decision (y_A). The appropriateness of such an assumption would depend on context, requiring a close appraisal of the circumstances in administrative law and regulation where an agency’s information is reliably encapsulated in its record. Although we do not work through details of this extension here, our core arguments are likely to carry over (with some caveats) to the case where A ’s signal is unobservable. In fact, if the median panelist and the agency share similar ideologies, our re-

³⁸Such a comparison may also bear on the issue of whether judicial ideologies exhibit drift over time (say, on the Supreme Court). Our model suggests that episodes of apparent drift could actually be due to changing information production patterns that coincide with changing ideological compositions of the Court.

sults tend to become even *more* pronounced. For example, suppose a Democrat agency is reviewed by either a DDD or DDR panel. The agency’s decision signals to the majority panelists that a politically-aligned actor observed a signal that they would likewise find persuasive, even if they cannot discern how strong that signal was. Unable to conduct a targeted audit of only those cases that are true “close calls,” the Democrats’ rational response might simply be to rubber stamp *all* of the agency’s decisions. A Republican minority panelist, in contrast, is more likely to retain an incentive to audit, but (just as above) she will do so only for A’s *liberal* pronouncements. Consequently, if A’s information were not observable, there can be equilibria where majority panelists never audit, and minority panelists (if any) categorically engage in one-sided auditing, reproducing (and even accentuating) the panel effects predicted in our baseline model.³⁹

Alternatively, one could perturb the informational environment at the *deliberation* stage, permitting auditing panelists to misrepresent (or selectively disclose) information to their colleagues. A panelist might, for example, misrepresent the *extensive* margin of her auditing efforts, covering up (perhaps at a cost) whether she has taken a hard look. Or, an auditing panelist might misrepresent the *intensive* margin of her efforts, falsifying or distorting (again, perhaps at some cost) the content of the signal she observed. It is relatively simple to extend our model to allow for misrepresentation on the extensive margin. So long as a judge’s ideological leaning is known (or accurately conjectured) by other panelists, it will be common knowledge whether she has an incentive to audit. The silence of a judge known to possess such an incentive creates an (accurate) inference by others that she discovered information inconsistent with her preferred position. In a manner akin to the unraveling phenomenon in information games (e.g., Milgrom and Roberts 1986), the equilibria identified above would substantially persist.

Misrepresentation on the intensive margin introduces a more complicated signaling game to our baseline model. We conjecture, however, that such an extension could entail similar effects. Here, non-auditing panelists, wary of falsification, would rationally interpret the content of the auditing judge’s signal in light of her ideology. When a fully separating equilibrium obtains, panelists can accurately decode the auditor’s signal, producing essentially the same equilibria described above. Under a complete pooling equilibrium, in contrast, the auditing judge sends uninformative signals, and other panelists simply ignore her. Anticipating this reaction, of course, the auditing judge would never collect the signal to begin with. This outcome would be identical to the baseline model where the cost of auditing (*c*) is prohibitive, and accordingly our model would not predict any panel effects. There may also be partially revealing equilibria, where some judges are willing to bear the cost of falsification, while others (those with less at stake) are not. In such equilibria, non-auditing panelists would selectively discount the auditor’s signal. We conjecture that in such equilibria,

³⁹Of course, the categorical nature of auditing in this case also implies that there can be contexts when *both* Democrats and Republicans audit, or when *neither* does. All told, our panel effects prediction would persist in the aggregate.

the severity of the panel’s discount increases as the auditing panelist’s ideology grows more extreme. Eventually, the marginal returns to diversity would dissipate for extremely ideological judges, who would lack credibility. Nevertheless, our core results would persist for judges falling inside this credibility threshold.

Yet a third extension of the model’s information structure would allow each panelist to draw a statistically independent auditing signal – either simultaneously with others or in sequence. Were this possible, *multiple judges* might choose to audit A ’s decision, each in an effort to sway the median voter. We have explored with this extension in a sequential setting, and it tends to produce a nested version of our baseline model, where extreme panelists engage in an iterative tug-of-war for the median voter’s favor. For example, suppose the median panelist initially leans liberal based on the agency’s developed facts. Under our baseline model, the most conservative panelist has the strongest incentive to audit. If he does, he may uncover information that pulls the median judge over to the conservative policy. In response, the most liberal panelist may *herself* rationally choose to take another draw, hoping to uncover information that wins back the median panelist. If she is successful, then yet another conservative judge may audit, and so forth down the line until the costs of the next draw are prohibitive. Viewed in this light, independent draws on v would likely amplify the deliberative dynamics that our baseline model exposes.

We might also extend the policy space beyond two distinct outcomes. For example, one could introduce a centrist policy option ($y = 0$) in addition to the liberal and conservative ones. This extension turns out to be relatively straightforward within our model, and has the effect of dampening *all judges*’ incentives to audit. For the median judge, a richer set of policy choices affords her the opportunity to fine tune the outcome to her ideal point and to her *a priori* information, which reduces both the costs of error and the value of additional information. Consequently, with more policy options it can become more attractive simply to remain uninformed and adopt the centrist position rather than to invest in additional information. The more ideological judges will also value additional information less, but they will still have incentives that are qualitatively similar to those in our benchmark model.⁴⁰

Another obvious – but possibly difficult – extension is to endogenize the Agency’s decision to do research. In our baseline model, the agency simply observes z ; there is no strategic choice involved. A literature going back at least to Gilligan and Krehbiel (1997),⁴¹ however, investigates the incentives of an administrative agency (or legislative committee) to gather information and expertise as a consequence of delegated authority. This literature has been extended to consider judicial oversight (Stephenson, 2007, 2008) and its effects on an Agency’s decision to gather expertise. A sophisticated court will tend to consider the feedback effects of its decision rule on the Agency’s decision, and will incorporate these effects into its rule of review. We could follow this

⁴⁰The principal difference is that with multiple policy choices, there may now be multiple disjoint auditing interval ranges around each of the median judge’s point of indifference between two ordered outcomes.

⁴¹See also Bueno de Mesquita and Stephenson (2007).

path with our model, by (say) permitting the agency to make a strategic investment in the precision of its initial signal, while anticipating how judicial review subsequently plays out.

Finally, we might attempt to embed our model within a multiple level auditing game. (See generally George (1999); George and Solomine (2001)). If we were to include the full circuit (for *en banc* review) and the Supreme Court, we would have four levels. Work is starting to be done with three levels, focusing on the full circuit’s decision to review.⁴² Indeed, Clark (2009) provides an elaborate empirical test of granting *en banc* review within a three-level principal agent framework, but does not provide a formal model. The equilibria of these models can be complex – a fact that may explain why some recent work (e.g., Landes and Posner 2009) fail to find panel effects at the Supreme Court level even though finding evidence at the circuit court level. Because our model provides a general framework for analyzing endogenous information production in arbitrarily sized panels, however, it may lend itself to such an extension.

6 Policy Implications

In our framework, mixed panels produce more information, which – through deliberation – brings about more informed decisions. While this seems intuitively desirable, it need not always follow that more informed decisions are always optimal, for at least two reasons: First, information in this model is purchased at a cost; even if majority panelists are eventually persuaded with new information, it does not imply that the added information was worth its cost from *social* perspective. Second, the additional information is generated instrumentally and is therefore likely skewed towards the interests of parties and political elites. If those interests do not coincide with the general interests of the citizenry, the additional information may not represent a real public good. These concerns aside, however, our analysis may lend at least *some* theoretical support to suggestions that we encourage (or even require) mixed panels within the federal judiciary (Schanzenbach and Tiller 2008).

Our framework does not directly allow us to make strong claims about legal doctrine, because doctrine is not a necessary ingredient of our model. However, it is a *possible* ingredient: as noted above: one plausible interpretation of our model is that the extra signal harvested by auditing panelists consists of undiscovered precedents, statutes, regulations or other persuasive authority over the issue. Under this interpretation, our model may suggest that mixed panels do a better job of uncovering and adhering to doctrine than do homogenous panels.⁴³

Our model may also have implications for the burgeoning theoretical and empirical literature on Supreme Court appointments. In this literature, the Senate and President observe the departure of a member of the Supreme Court, and then bargain in some structural setting over the new appointee (e.g., Kre-

⁴²E.g., Kastellic (2010). Revesz (1997) at pg. 1747, investigates a “hierarchical constraint” hypothesis that stems from the possibility of Supreme Court review.

⁴³On the other hand, our model can say little about writing opinions (majority or dissent).

hbiel 2007). Both the Senate and the President evaluate new appointees by referring to their expected votes. In turn, these expectations are conventionally thought to be the function of each potential nominee’s individual characteristics. Our model (and the empirical literature that attracted us to it), however, suggests that the voting proclivities of a new appointee may be significantly more complex than this and turn on who is empaneled with him/her. Moreover, the new appointee may also perturb the voting proclivities of incumbent members of the court. Embedding this feature into the appointment literature is an interesting (and in our mind worthwhile) challenge.

From a topical perspective, there may also be a number of applications of our approach. For example, many of the information production / deliberation intuitions analyzed above carry over to other multimember political decision makers, such as administrative committees or agencies themselves. Our approach may also dovetail with and contribute to the literature about the endogenous formation of peer groups through *homophily* (i.e., connection and information sharing among philosophically-allied individuals⁴⁴). Within organizational theory, our analysis may shed light on the extent to which heterogeneity of world views among block shareholders or corporate board members may better inform corporate decisions.⁴⁵ Similarly, our approach may shed light on the conditions under which having single versus numerous large block shareholders in the ownership structure of a company can facilitate efficient endogenous information production – a question that has become increasingly important of late.⁴⁶

7 Conclusion

In this paper, we have presented a simple information-based model of panel deliberation at the circuit court level. Proposition 1 (and associated corollaries) captures central insights from the panel effects literature. Specifically, we have illustrated how mixed panels may induce equilibria manifesting the markers of moderation among majority (and even minority) panelists. The type of moderation we predict is not an artifact of endogenous preferences, collegiality, group cohesion or whistleblowing *per se*, but rather the product of endogenous patterns of information production, developed and provided by panelists who have diverse ideological commitments. In at least some respects, our argument is consistent with the claim that mixed panels produce not only different results, but also better results than their homogenous counterparts. Our information-based account joins a group of theories attempting to explain the phenomena of both majority and minority judges as a function of panel composition, and fu-

⁴⁴See, e.g., Currarini et al. (2009).

⁴⁵For example, the now well-documented disagreements between Patricia Dunn (a governance-oriented director) and Tom Perkins and Jay Keyworth (two strategy-oriented directors) on the Hewlett Packard Board may have some benefits even while it potentially foments internal conflict. See, e.g., Anders & Murray (2006).

⁴⁶See, e.g., Yucaipa American Alliance Fund LLC v. Reggio et al, C.A. No. 5465-VCS, (Del. Ch. 2010) (challenging a poison pill that grandfathered in a pre-existing 37-percent shareholder while being triggered by all others who surpassed 15 percent).

ture empirical tests must sort out which theory has greatest explanatory power in practice. We have suggested a few promising routes for such tests, but we leave their execution for another day.

8 References

1. Anders, G. and Alan Murray. 2006. "Boardroom Duel: Behind H-P Chairman's Fall, Clash with a Powerful Director; the Cautious Patricia Dunn and Flashy Tom Perkins were a Combustible Pair; Overcoming a 'Respect Gap' " *The Wall Street Journal Wall St. Journal* (10/9/06).
2. Asch, Solomon E. 1955. "Opinions and social pressure," 193 *Scientific American* 31-35.
3. Atkins, Burton M. 1973. "Judicial Behavior and Tendencies Towards Conformity in a Three Member Small Group: A Case Study of Dissent Behavior on the U.S. Court of Appeals," 54 *Social Science Quarterly* 41.
4. Atkins, Burton M. and Justin J. Green. 1976. "Consensus on the United States Courts of Appeals: Illusion or Reality?" 20 *American Journal of Political Science* no. 4: 735-748.
5. Aumann, Robert J. 1976. "Agreeing to Disagree," 4 *Ann. Stat.* 1236-1239.
6. Bawn, Kathleen. 1997. "Choosing Strategies To Control the Bureaucracy: Statutory Constraints, Oversight, and the Committee System," 13 *Journal of Law, Economics & Organization* 101-126.
7. Boyd, Christina L., Lee Epstein, and Andrew D. Martin. 2011. "Untangling the Causal Effects of Sex on Judging," *American Journal of Political Science* (forthcoming).
8. Bressman, Lisa Schultz. 2007. "Procedures as Politics in Administrative Law," 107 *Columbia Law Review* 1749-1821.
9. Bueno de Mesquita, Ethan and Matthew C. Stephenson. 2007. "Regulatory Quality Under Imperfect Oversight," 101 *American Political Science Review* no. 3: 605-620.
10. Cameron, Charles M., Jeffrey A. Segal, and Donald R. Songer. 2000. "Strategic Auditing in a Political Hierarchy: An Informational Model of the Supreme Court's Certiorari Decisions," 94 *American Political Science Review* 101-16.
11. Cameron, Charles M., and Lewis A. Kornhauser. 2006. "Appeals Mechanisms, Litigant Selection, and the Structure of Judicial Hierarchies," in J. Rogers, R. Flemming, and J. Bond, eds., *Institutional Games and the U.S. Supreme Court*. Charlottesville, Virginia: University of Virginia Press. 173-205.

12. Clark, Tom S. 2009. "A Principal-Agent Theory of En Banc Review," 25 *Journal of Law, Economics & Organization* 55-79.
13. Cohen, Linda R. and Matthew L. Spitzer. 1996. "Judicial Deference to Agency Action: A Rational Choice Theory and an Empirical Test," 69 *Southern California Law Review* 431-476.
14. — — —. 1994. "Solving the Chevron Puzzle," 57 *Law and Contemporary Problems* 65-110.
15. Cox, Adam B. and Thomas J. Miles. 2007. "Judging the Voting Rights Act" 108 *Columbia Law Review* 1-54.
16. Cross, Frank B. 2007. *Decision-making in the U.S. Courts of Appeals*. Palo Alto: Stanford University Press.
17. Cross, Frank B. and Emerson H. Tiller. 2008. "Judicial Partisanship and Obedience to Legal Doctrine: Whistleblowing on the Federal Courts of Appeal," 107 *Yale L. J.* 2155-76.
18. — — —. "Whistleblowing on the Federal Courts of Appeal," 107 *Yale Law Journal* 2155-2176 (May 1998).
19. Currarini, Sergio, Matthew O. Jackson and Paolo Pin. 2009. "An Economic Model of Friendship: Homophily, Minorities, and Segregation," 77 *Econometrica* 1003-45.
20. Czarnezki, Jason P. 2008. "An Empirical Investigation of Judicial Decisionmaking, Statutory Interpretation, and the Chevron Doctrine in Environmental Law," 79 *University Colorado Law Review* 767-823.
21. DeFigueiredo, John M., and Charles M. Cameron. 2008. "Endogenous Cost Lobbying: Theory and Evidence," UCLA Business School Working Paper.
22. DeGroot, Morris. 2004. *Optimal Statistical Decisions* (Wiley & Sons, 2nd Ed).
23. Dewatripont, M. and J. Tirole. 1999. "Advocates," 107 *Journal of Political Economy* 1-39.
24. Edwards, Harry T. and Michael A. Livermore. 2009. "Pitfalls of Empirical Studies that Attempt to Understand the Factors Affecting Appellate Decision-making" 58 *Duke Law Journal* 1895-89.
25. Epstein, David and Sharyn O'Halloran. 1999. "Delegating Powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers," Cambridge University Press.

26. Epstein, Lee, William Landes and Richard Posner. 2011. "Why (And When) Judges Dissent: A Theoretical and Empirical Analysis," 1 *Journal of Legal Analysis* 101-37.
27. Epstein, Lee, Andrew D. Martin, Jeffery A. Segal, and Chad Westerland. 2007. "The Judicial Common Space," 23 *Journal of Law, Economics and Organization* 303-25.
28. Epstein, Lee, Andrew D. Martin, Kevin M. Quinn and Jeffrey A. Segal. 2007. "Ideological Drift among Supreme Court Justices: Who, When, and How Important?" 101 *Northwestern University Law Review* 1483-1542.
29. Epstein, Lee, and Jeffrey A. Segal. 2005. *Advice and Consent: The Politics of Federal Judicial Appointments*. New York, New York: Oxford University Press.
30. Eskridge, Jr., William N. and John Ferejohn. 1992. "The Article I, Section 7 Game," 80 *Georgetown Law Journal* 523-564.
31. Farhang, Sean, and Gregory Wawro. 2004. "Institutional dynamics on the U.S. Court of Appeals: Minority representation under panel decision making," 20 *Journal of Law, Economics, & Organization* 299-330.
32. Fischman, Joshua B. 2009. "Estimating Preferences of Appellate Judges: A Model of 'Consensus Voting'," Working Paper. Available at SSRN: <http://ssrn.com/abstract=1361348>.
33. George, Tracey E. 1999. "The Dynamics and Determinants of the Decision to Grant En Banc Review," 74 *Washington Law Review* 213-74.
34. George, Tracey E., and Michael E. Solimine. 2001. "Supreme Court Monitoring of the United States Courts of Appeals En Banc," 9 *Supreme Court Economic Review* 171-204.
35. Gilligan, Thomas W, and Keith Krehbiel. 1997. "Specialization Decisions Within Committee," 13 *Journal of Law, Economics & Organization* 366-386.
36. Goldman, Sheldon. 1968. "Conflict and Consensus in the United States Courts of Appeals," *Wisconsin Law Review* 461-80.
37. Gentzkow, Matthew and Jesse M. Shapiro. 2006. "Media Bias and Reputation" 114 *Journal of Political Economy* 280-316.
38. Green, Justin J. 1986. Parameters of Dissensus on Shifting Small Groups, in S. Goldman and C.M. Lamb eds., *Judicial Conflict and Consensus: Behavioral Studies of American Appellate Courts*, Lexington, Kentucky: University of Kentucky Press.

39. Hettinger, Virginia A., Stefanie A. Lindquist and Wendy L. Martinek. 2007. *Judging on a Collegial Court: Influences on Federal Appellate Decision Making (Constitutionalism and Democracy)*. Charlottesville: University of Virginia Press.
40. Kastellec, Jonathan P. 2007. "Panel Composition and Judicial Compliance on the US Courts of Appeals," 23 *Journal of Law, Economics and Organization* 421-441.
41. — — —. 2008, Panel Composition and Voting on the U.S. Courts of Appeals Over Time. Working paper, Princeton University.
42. — — —. 2010. Hierarchical and Collegial Politics on the U.S. Courts of Appeals. Working paper, Princeton University.
43. — — —. 2011. Racial Diversity and Judicial Influence on Appellate Courts. Working paper, Princeton University.
44. Kornhauser, Lewis A. 1992. "Modeling Collegial Courts II: Legal Doctrine," 8 *Journal of Law, Economics, and Organization* no.3 441-470.
45. — — —. 1992. "Modeling Collegial Courts I: Path Dependence," 12 *International Review of Law and Economics* 169-85.
46. Krehbiel, Keith. 2007. "Supreme Court Appointments as a Move-the-Median Game," 51 *American Journal of Political Science* no. 2 231-40.
47. Landes, William M. and Richard A. Posner. 2009. "Rational Judicial Behavior: A Statistical Study," 1 *The Journal of Legal Analysis* 775-831.
48. Landsman, Zinoviy and Emiliano A. Valdez. 2005. "Tail Conditional Expectations for Exponential Dispersion Models," 35 *Astin Bulletin* 189-210.
49. McNollGast. 1999. 'The Political Origins of the Administrative Procedure Act," 15 *Journal of Law, Economics and Organization* 180-217.
50. — — —. 1995. "Politics and the Courts: A Positive Theory of Judicial Doctrine and the Rule of Law," 68 *Southern California Law Review* 1631-89.
51. Maveety, Nancy. 2005. "The Study of Judicial Behavior and the Discipline of Political Science" in Nancy Maveety, ed., *The Pioneers of Judicial Behavior*. Ann Arbor, MI: University of Michigan Press. 1-51.
52. Miles, Thomas J. and Cass R. Sunstein. 2006. "Do Judges Make Regulatory Policy? An Empirical Investigation of Chevron," 73 *University of Chicago Law Review* 823-881.
53. — — —. 2008. 'The Real World of Arbitrariness Review," 75 *University of Chicago Law Review* 761-814.

54. Milgrom, Paul and John Roberts. 1986. "Relying on the Information of Interested Parties," 17 *Rand Journal of Economics* 18-32.
55. Moraski, Bryon and Charles Shipan. 1999. "The Politics of Supreme Court Nominations: A Theory of Institutional Constraints and Choices," 43 *American Journal of Political Science* 1069-1095.
56. Myers, David G. 1975. "Decision-Induced Attitude Polarization," 28 *Human Relations* 699-714.
57. Peresie, Jennifer L. 2005. "Female Judges Matter: Gender and Collegial Decision-making in the Federal Appellate Courts." 114 *Yale Law Journal* 1759-90.
58. Posner, Richard A. 2002. *How Do Judges Think?* Cambridge, Massachusetts: Harvard University Press.
59. Revesz, Richard L. 1997. "Environmental Regulation, Ideology, and the D.C. Circuit," 83 *Virginia Law Review* 1717-72.
60. Rohde, David W. and Kenneth A. Shepsle. 2007. "Advising and Consenting in the 60-Vote Senate: Strategic Appointments to the Supreme Court," 69 *Journal of Politics* 664-77.
61. Rodriguez, Daniel B. and Barry Weingast. 2007. "The Paradox of Expansionist Statutory Constructions," 101 *Northwestern Law Review* 1207-55.
62. Schanzenbach, Max M. and Emerson H. Tiller. 2007. "Strategic Judging Under the US Sentencing Guidelines: Positive Political Theory and Evidence," 23 *Journal of Law, Economics and Organization* 24-56.
63. — — —. 2008. "Reviewing the Sentencing Guidelines: Judicial Politics, Empirical Evidence, and Reform," 75 *University of Chicago Law Review* 715-760.
64. Segal, Jeffrey A., and Harold J. Spaeth. 2002. *The Supreme Court and the Attitudinal Model Revisited*. New York: Cambridge University Press.
65. Shavell, Steven. 2009. "On the Design of the Appeals Process: The Optimal Use of Discretionary Review versus Direct Appeal," *Harvard Law School Discussion Paper* No. 625.
66. Sisk, Gregory C. and Michael Heise. 2005. "Judges and Ideology: Public and Academic Debates About Statistical Measures," 99 *Northwestern University Law Review* 743-804.
67. Smith, Joseph L. 2007. "Presidents, Justices, and Deference to Administrative Action," 23 *Journal of Law, Economics and Organization* 346-64.

68. Songer, Donald R. 1982. "Consensual and Nonconsensual Decisions in Unanimous Opinions of the United States Courts of Appeals," 26 *American Journal of Political Science* 225-239.
69. Songer, Donald, Jeffrey A. Segal, and Charles M. Cameron. 1994. "The Hierarchy of Justice: Testing a Principal-Agent Model of Supreme Court-Circuit Court Interactions," 36 *American Journal of Political Science* 673-96.
70. Spitzer, Matthew L. and Eric Talley. 2000. "Judicial Auditing," 29 *Journal of Legal Studies* 649-683.
71. Stephenson, Matthew C. 2006. "The Strategic Substitution Effect: Textual Plausibility, Procedural Formality, and Judicial Review of Agency Statutory Interpretations," 120 *Harvard Law Review* 528-572.
72. ——. 2008. "Evidentiary Standards and Information Acquisition in Public Law," 10 *American Law and Economics Review* 351-387.
73. ——. 2007. "Bureaucratic Decision Costs and Endogenous Agency Expertise," 23 *Journal of Law, Economics and Organization* 469-98.
74. ——. 2009. "Statutory Interpretation by Administrative Agencies," in D. Farber & A. O'Connell, eds., *Research Handbook in Public Law and Public Choice*. Northampton, MA: Edward Elgar Publishing. 285-332.
75. Strauss, Peter L. 1987. "One Hundred Fifty Cases Per Year: Some Implications of the Supreme Court's Limited Resources for Judicial Review of Agency Action," 87 *Columbia Law Review* 1093-1136.
76. Sunstein, Cass R., David Schkade, Lisa Ellman and Andres Sawicki. 2006. *Are Judges Political? An Empirical Analysis of the Federal Judiciary*. Washington, DC: Brookings Institution Press.
77. Sunstein, Cass, David Schkade, and Lisa Ellman. 2004. "Ideological Voting on Federal Courts of Appeals: A Preliminary Investigation," 90 *Virginia Law Review*, 301-54.
78. Tacha, Deanell Reece. 1995. "The 'C' Word: On Collegiality," 56 *Ohio State Law Journal* 585-92.
79. Talley, Eric. 1999. "Precedential Cascades: An Appraisal," 73 *Southern California Law Review* 87-137.
80. Tiller, Emerson H. and Frank B. Cross. 1999. "A Modest Proposal for Improving American Justice," 99 *Columbia Law Review* 215-34.
81. Wald, Patricia M. 1999. "A Response to Tiller and Cross," 99 *Columbia Law Review* 235-61.

9 Appendix

This appendix includes some basic identities and mathematical derivations that enter into the analysis, as well as proof of core propositions.

9.1 Distributional Identities

Analyzing the model in the text requires some manipulation of certain probability distributions. For the reader's reference, some of the key identities are stated here. Recall from the model that the "true" state of the world, X , is distributed $N(\mu, \frac{1}{\tau})$; Player A's signal Z is distributed such that $(Z|X) \sim N(x, \frac{1}{\gamma})$; and Player J's signal V is distributed such that $(V|X) \sim N(x, \frac{1}{\rho})$. Because each of X , Z , and V is distributed normally, their various conditional random variables remain normal as well (see DeGroot 2004). Specifically, Table A1 reports the mean and precision (the inverse of the variance) for five conditional random variables of interest:

Cond. RV	Distribution	Mean	Precision
$X Z, V$	<i>Normal</i>	$\frac{\tau\mu + \gamma z + \rho v}{\tau + \gamma + \rho}$	$\tau + \gamma + \rho$
$X Z$	<i>Normal</i>	$\frac{\tau\mu + \gamma z}{\tau + \gamma}$	$\tau + \gamma$
$X V$	<i>Normal</i>	$\frac{\tau\mu + \rho v}{\tau + \rho}$	$\tau + \rho$
$V Z$	<i>Normal</i>	$\frac{\tau\mu + \gamma z}{\tau + \gamma}$	$\frac{(\tau + \gamma)\rho}{\tau + \gamma + \rho}$
$Z V$	<i>Normal</i>	$\frac{\tau\mu + \rho v}{\tau + \rho}$	$\frac{(\tau + \rho)\gamma}{\tau + \gamma + \rho}$

Table A1. Conditional Distributions

In addition, our model makes use of the (so-called) *Tail Conditional Expectation*. Specifically, if $X \sim N(\alpha, \frac{1}{\beta})$, then the expectation of X conditional on $X \geq \hat{x}$ is given by:

$$E(X|X \geq \hat{x}) = \alpha + \frac{1}{\sqrt{\beta}} \cdot \left(\frac{\phi((\hat{x} - \alpha)\sqrt{\beta})}{1 - \Phi((\hat{x} - \alpha)\sqrt{\beta})} \right) \quad (\text{A1})$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ represent the standard normal probability density and cumulative distribution functions, respectively (Landsman and Valdez 2005).

Finally, our model requires identifying various order statistics on the set of panelist ideologies, Θ . Consider a vector of realizations $\Theta \equiv \{\theta_1, \theta_2, \dots, \theta_{2M-1}\}$, drawn independently from an identical distribution $H(\theta)$ with associated density $h(\theta)$. Without loss of generality, we can reorder Θ in terms of order statistics $\{\theta_{(1)}, \dots, \theta_{(M)}, \dots, \theta_{(2M-1)}\}$. Define $\hat{\Theta} \subset \Theta$ as the 3-tuple $\{\theta_{(1)}, \theta_{(M)}, \theta_{(2M-1)}\}$, representing the minimal, median, and maximal elements of Θ . The k -th order statistic, or $\theta_{(k)}$, has a probability density function given by:

$$h_{(k)}(\theta_{(k)}) = \frac{(2M-1)!}{(k-1)!(2M-1-k)!} (H(\theta_{(k)}))^{k-1} (1 - H(\theta_{(k)}))^{(2M-1-k)} h(x) \quad (\text{A2})$$

Applying this expression iteratively, the joint pdf of $\hat{\Theta}$, in terms of $H(\theta)$ and $h(\theta)$, is as follows:

$$\begin{aligned} f(\hat{\Theta}) &= \frac{(2M-1)!}{(M-2)!^2} \cdot h(\theta_{(1)}) \cdot h(\theta_{(M)}) \cdot h(\theta_{(2M-1)}) \\ &\quad \times (H(\theta_{(M)}) - H(\theta_{(1)}))^{(M-2)} (H(\theta_{(2M-1)}) - H(\theta_{(M)}))^{(M-2)} \end{aligned} \quad (\text{A3})$$

9.2 Derivation of Expected Payoff for Uninformed Judge

Consider a judge with ideology θ_i sitting on an uninformed panel with ideological profile Θ . Judge i 's expected payoff if informed (conditional on z) is given by:

$$\begin{aligned} \pi_U(\theta|z, \theta_{(M)}) &= -E_{x|z} \left\{ ((x + \theta_i) - y)^2 | z \right\} \\ &= \begin{cases} -E_{x|z} \left\{ (x + \theta_i + 1)^2 | z \right\} & \text{if } z \leq z_M^U \\ -E_{x|z} \left\{ (x + \theta_i - 1)^2 | z \right\} & \text{else} \end{cases} \\ &= -E_{x|z} \left(x^2 + 2x(\theta_i + 1) + (\theta_i + 1)^2 \right) + 4 \begin{cases} 0 & \text{if } z \leq z_M^U \\ E_{x|z} \{ (\theta_i + x) | z \} & \text{else} \end{cases} \\ &= \left(\frac{1}{\tau + \gamma} + \left(\frac{\tau\mu + \gamma z}{\tau + \gamma} + (\theta_i + 1) \right)^2 \right) + \begin{cases} 0 & \text{if } z \leq z_M^U \\ 4 \left(\theta_i + \frac{\tau\mu + \gamma z}{\tau + \gamma} \right) & \text{else} \end{cases}, \end{aligned} \quad (\text{A4})$$

which is the expression given in the text.

9.3 Derivation of Expected Payoff for Informed Judge

Assuming the panel becomes informed of v , it will issue a decision y_M^I , consistent with the median judge's ideology $\theta_{(M)}$. Expected payoff of any judge on the panel

with ideology θ_i is thus:

$$\begin{aligned}
\pi_I(\theta_i|z, \theta_{(M)}) &= -E_{v|z} \left\{ E_{x|z,v} (x + \theta_i - y_M^I)^2 | z, v \right\} \\
&= E_{v|z} \left[\begin{cases} -E_{x|v,z} \left\{ (x + \theta_i + 1)^2 | z, v \right\} & \text{if } v \leq v_M^I \\ -E_{x|v,z} \left\{ (x + \theta_i - 1)^2 | z, v \right\} & \text{if } v > v_M^I \end{cases} \middle| z, v \right] \\
&= - \left(\frac{1}{\tau + \gamma} + \left(\frac{z\gamma + \tau\mu}{\tau + \gamma} + (\theta_i + 1) \right)^2 \right) + E_{v|z} \left\{ \begin{matrix} 0 & \text{if } v \leq v_M^I \\ 4 \left(\frac{\tau\mu + \gamma z + \rho v}{\tau + \gamma + \rho} + \theta_i \right) & \text{if } v > v_M^I \end{matrix} \right\} \\
&= - \left(\frac{1}{\tau + \gamma} + \left(\frac{z\gamma + \tau\mu}{\tau + \gamma} + (\theta_i + 1) \right)^2 \right) \\
&\quad + 4 \cdot \left(\theta_i + \left(\frac{\tau\mu + \gamma z}{\tau + \gamma} \right) \right) \left(1 - \Phi \left(- \frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\rho + \tau + \gamma)(\tau + \gamma)}}} \right) \right) \\
&\quad + 4 \cdot \sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}} \cdot \phi \left(- \frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\rho + \tau + \gamma)(\tau + \gamma)}}} \right)
\end{aligned} \tag{A5}$$

9.4 Proofs of Lemmas 1-5

Lemma 1: For the median judge, $\Delta(\theta_{(M)}|z, \theta_{(M)})$ is maximal at $z = z_M^U$, and falls symmetrically in both directions as z moves away from z_M^U . Consequently, when panel ideologies are homogenous, the auditing range also will constitute a symmetric interval around z_M^U .

Proof: First, note that $\left(\theta_i + \frac{z\gamma + \tau\mu}{\tau + \gamma} \right) \Big|_{z=z_M^U} = 0$. Therefore, $\Delta(\theta_{(M)}|z, \theta_{(M)}) \Big|_{z=z_M^U}$ simplifies to:

$$\begin{aligned}
\Delta(\theta_{(M)}|z, \theta_{(M)}) &= 4 \cdot \sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}} \cdot \phi(0) \\
&= 4 \cdot \phi(0) \cdot \sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}} \\
&= \sqrt{\frac{8}{\pi}} \cdot \left(\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)} \right)
\end{aligned} \tag{A6}$$

Note further that the standard normal density $\phi(x)$ is maximized at $x = 0$, and thus the first additive term of (4) is maximized when $z = z_M^U$. As to the second term of (4), it is easily verified that its value is negative for all values of $z \neq z_M^U$. Thus, since both additive terms of (4) are maximized at z_M^U , so must their sum. The symmetry of $\Delta(\theta_{(M)}|z, \theta_{(M)})$ around $z = z_M^U$ follows immediately from the symmetry of the standard normal distribution around 0. *QED*

Lemma 2: If judge i is more conservative than the median judge, so that $\theta_i > \theta_{(M)}$:

- Judge i values information more than the median judge when $z \leq z_M^U$ and less than the median judge when $z > z_M^U$.
- The extent to which the more conservative judge's valuation exceeds / falls short of the median judge's increases strictly in θ_i .

If judge i more liberal than the median judge, so that $\theta_i < \theta_{(M)}$:

- Judge i values information more than the median judge when $z \geq z_M^U$ and less than the median judge when $z < z_M^U$.
- The extent to which the more liberal judge's valuation exceeds / falls short of the median judge's decreases strictly in θ_i .

Proof: An equivalent way to express the value of information for the non-median judge is to consider the degree to which judge i 's valuation of auditing exceeds that of the median judge. Denoting this valuation gap as $\xi(\theta_i, \theta_{(M)}, z)$, the following expression emerges:

$$\begin{aligned} \xi(\theta_i, \theta_{(M)}, z) &= \Delta(\theta_i | z, \theta_{(M)}) - \Delta(\theta_{(M)} | z, \theta_{(M)}) \\ &= 4 \cdot (\theta_i - \theta_{(M)}) \cdot \begin{cases} 1 - \Phi\left(-\frac{(\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma})}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) & \text{if } z \leq z_M^U \\ -\Phi\left(-\frac{(\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma})}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) & \text{if } z > z_M^U \end{cases} \end{aligned} \quad (\text{A7})$$

The statements in the Lemma come directly from inspection and/or differentiation of $\xi(\theta_i, \theta_{(M)}, z)$. *QED*

Lemma 3: When $z < z_M^U$ the most conservative judge (with ideology $\theta_{(2M-1)}$) has the maximal incentive of all panelists to audit. Similarly, when $z > z_M^U$, the most liberal judge (with ideology $\theta_{(1)}$) has the maximal incentive to audit. If $z = z_M^U$, the most conservative (most liberal) panelist has the greatest incentive to audit when $(\theta_{(2M-1)} - \theta_{(M)})$ is larger (smaller) than $(\theta_{(M)} - \theta_{(1)})$.

Proof: Direct implication of Lemma 2.

Lemma 4: If $c \leq c(\hat{\Theta}, z)$, the panel will audit (and thus learn v) where

$$\begin{aligned} c(\hat{\Theta}, z) &= 4\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}} \cdot \phi\left(-\frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) \\ &\quad + \begin{cases} 4\left(\theta_{(2M-1)} + \frac{\tau\mu + \gamma z}{\tau + \gamma}\right) \cdot \left(1 - \Phi\left(-\frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right)\right) & \text{if } z \leq z_M^U \\ -4\left(\theta_{(1)} + \frac{\tau\mu + \gamma z}{\tau + \gamma}\right) \cdot \Phi\left(-\frac{\theta_{(M)} + \frac{z\gamma + \tau\mu}{\tau + \gamma}}{\sqrt{\frac{\rho}{(\tau + \gamma + \rho)(\tau + \gamma)}}}\right) & \text{if } z > z_M^U \end{cases} \end{aligned} \quad (\text{A8})$$

This criterion implicitly defines strictly positive (but possibly asymmetric) auditing interval $[\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})]$ around z_M^U .

Proof: Direct implication of Lemmas 2-3.

Lemma 5: *Given the equilibrium behavior of a panel with configuration $\hat{\Theta}$, A will make the conservative decision if and only if:*

$$4(1-\pi) \left(\frac{\tau\mu + \gamma z}{\tau + \gamma} + \theta_A \right) \geq \pi \varepsilon \cdot \left(\begin{array}{l} E \left(\Pr \left\{ z < \underline{z}(\hat{\Theta}) \right\} | z \right) - E \left(\Pr \left\{ z > \bar{z}(\hat{\Theta}) \right\} | z \right) \\ + E \Pr \left\{ z \in \left[\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta}) \right] \cap v < v_M^I | z \right\} \\ - E \Pr \left\{ z \in \left[\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta}) \right] \cap v \geq v_M^I | z \right\} \end{array} \right)$$

Proof: Conditional on knowing z , if A issues the conservative decision ($y_A = 1$), her expected payoff will be:

$$\begin{aligned} & - (1-\pi) \cdot E \left((x + \theta_A - 1)^2 | z \right) \tag{A9} \\ & - \pi \varepsilon \cdot \left(\int \int \int_{\hat{\Theta} | z < \underline{z}(\hat{\Theta})} f(\hat{\Theta}) d\hat{\Theta} + \int \int \int_{\hat{\Theta} | z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})]} \Phi(-(\theta(\tau + \gamma) + z\gamma + \tau\mu)) f(\hat{\Theta}) d\hat{\Theta} \right) \end{aligned}$$

where $f(\hat{\Theta})$ is as derived above. If A issues the liberal decision ($y_A = -1$), in contrast, her expected payoff will be:

$$\begin{aligned} & - (1-\pi) \cdot E \left((x + \theta_A + 1)^2 | z \right) \tag{A10} \\ & - \pi \varepsilon \cdot \left(\int \int \int_{\hat{\Theta} | z > \bar{z}(\hat{\Theta})} f(\hat{\Theta}) d\hat{\Theta} + \int \int \int_{\hat{\Theta} | z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})]} (1 - \Phi(-(\theta(\tau + \gamma) + z\gamma + \tau\mu))) f(\hat{\Theta}) d\hat{\Theta} \right) \end{aligned}$$

Consequently, A will make the conservative decision if and only if the difference between these two expressions is positive, or:

$$4(1-\pi) \left(\frac{\tau\mu + \gamma z}{\tau + \gamma} + \theta_A \right) + \pi \varepsilon \cdot \left(\begin{array}{l} \int \int \int_{\hat{\Theta} | z > \bar{z}(\hat{\Theta})} f(\hat{\Theta}) d\hat{\Theta} - \int \int \int_{\hat{\Theta} | z < \underline{z}(\hat{\Theta})} f(\hat{\Theta}) d\hat{\Theta} \\ + \int \int \int_{\hat{\Theta} | z \in [\underline{z}(\hat{\Theta}), \bar{z}(\hat{\Theta})]} (1 - 2\Phi(-(\theta(\tau + \gamma) + z\gamma + \tau\mu))) f(\hat{\Theta}) d\hat{\Theta} \end{array} \right) \geq 0 \tag{A11}$$

Rearranging and substituting appropriate expectation operators yields the expression in the text. *QED.*