Biased Judgments without Biased Judges:
How Legal Institutions Cause Errors*

Ryan Hübert[†]

University of California, Davis

November 5, 2019

Abstract

Certain features of the legal system, such as litigant-driven appeals and random assignment of judges to cases, are supposed to prevent biased and error-prone decision-making by judges. I analyze a formal model illustrating this is not always true, even when judges are personally unbiased. First, an unbiased trial judge may make systematically biased judgments when her decisions can be appealed by a losing litigant. If the judge is sufficiently reversal averse, this bias favors higher resource litigants. Second, litigant-driven appeals allow unbiased judges to put less effort into resolving cases correctly. Third, when unbiased judges are assigned randomly to cases, they will produce more errors than a case assignment system allowing trial judges more freedom to select cases. The results confirm one contention of "team models" of judicial hierarchy—that litigant behavior is important—while also demonstrating the limits of another—that judicial hierarchies minimize adjudication errors.

Keywords: judicial politics, judicial hierarchy, trial courts, team models, formal theory

Supplementary material for this paper is available in the Online Appendix.

Officials in the federal judiciary regularly push back on claims that judges are biased. In late 2018, Chief Justice John Roberts issued a rare statement in response to President Trump's criticism of an "Obama judge" who had ruled against the administration. The Chief Justice insisted: "We do not have Obama judges or Trump judges, Bush judges or Clinton judges [...] What we have is an extraordinary group of dedicated judges doing their level best to do equal right to those appearing before them" (quoted in Sherman 2018). Some judges have also criticized academic research demonstrating political ideology affects judges' decisions. In response to Revesz (1997), which shows that Republican- and Democratic-appointed judges on the D.C. Circuit rule differently in EPA cases, D.C. Circuit Judge Harry T. Edwards argued that the findings "may mislead the unsuspecting [...] into thinking that judges are lawless in their decision making, influenced more by personal ideology than legal principles" (Edwards 1998, p. 1337).

The debate over judicial bias is rooted in a broader normative concern that judges—who are more independent and less accountable than other public officials—may systematically misapply the law to the specific cases they hear. In spite of frequent (and vehement) claims that judges are hard-working and unbiased, a set of institutions has developed, at least in part, to guard against this. On the front end of a case, judges are typically randomly assigned in order to prevent them from deciding cases in which they may be biased in favor or against a party. On the back end of a case, parties who lose in a lower court are generally afforded an opportunity to appeal to a higher court and obtain a reversal of a bad decision. Many of these institutions are designed to minimize the influence of individual judges over legal outcomes. The presumption is that, by doing so, these institutions will reduce the number of cases that will be resolved erroneously.

But is this presumption always correct? I present a formal model that illustrates how institutions that are meant to ensure the legal system correctly decides cases can perversely generate biased outcomes. Strikingly, this occurs even though I assume all judges in the model are completely unbiased. The model provides three core insights. First, an unbiased trial judge may make systematically biased judgments when subject to litigant-driven appellate review. This occurs be-

cause of their strategic anticipation of litigant appeals. As a result, even when staffed entirely by judges who share the same preferences over case outcomes, the judicial hierarchy can actually *encourage* errors. Second, litigant-driven appeals sometimes incentivize trial judges to lower the amount of effort they put into managing cases. Third, when a heterogeneous set of unbiased judges are assigned randomly to cases, they may produce more errors than under a case assignment system that allows trial judges freedom to select the cases they hear. Moreover, this has implications for judicial selection: diversity on the bench may promote high quality decision-making under flexible case assignment procedures, but not under random case assignment.

Due to the incentives that can emerge under appellate review, the analysis in this paper casts doubt on the conventional wisdom that unbiased judges will generally yield unbiased outcomes. It also helps clarify the scope conditions of existing theories of the judicial hierarchy that have shown how a "team" of unbiased judges sitting in a hierarchy can minimize the number of wrongly decided cases (e.g., Cameron and Kornhauser 2005 and 2006, from here forward CK 2005 and CK 2006). As I show in this paper, the judicial hierarchy in these team models only minimizes errors under certain conditions that may not always hold. In particular, those findings require either symmetry between litigants (as in CK 2006) or judges' indifference about being reversed (as in CK 2005). Relaxing both of these assumptions yields a qualitatively and normatively different understanding of how judges operate in a judicial hierarchy. Moreover, this has implications for how cases get assigned to judges since random assignment does not mitigate the potentially negative consequences of litigant-driven appellate review.

Law Application in the Judicial Hierarchy

The bulk of the time, energy and effort put into legal cases happens at the first stages of adjudication—in the trial courts. For example, in the year ending March 31, 2018, there were 346,608 civil and criminal cases terminated in the federal trial courts, as opposed to 51,832 in the lower federal ap-

pellate courts (Administrative Office of the United States Courts 2018, tables B–D). In addition to being a much larger presence in the judiciary's docket than appellate courts, trial courts also serve a unique purpose. They do not generally make new law, and they are only tangentially involved in the creation of doctrine. They are instead tasked with developing a deep understanding of individual cases so that they can figure out how to apply the law in those cases. This is often referred to as "law application." A key question is whether they do this well. In other words, if we take existing law as given, how many errors do courts make in resolving cases? These errors have two sources: mistakes and willful non-compliance.

On both counts, recent research presents a relatively optimistic view. Scholars have found little empirical evidence of non-compliance by trial judges (e.g., Boyd and Spriggs 2009; Boyd 2015; Epstein, Landes, and Posner 2013; Randazzo 2008; Schanzenbach and Tiller 2007). These studies largely demonstrate that trial judges strategically alter their decision-making to conform to their appellate court overseers. While compelling, this account sometimes under-emphasizes the role that litigants play in shaping judges' incentives. Indeed, a trial judge's decision-making does not just occur in the shadow of an appellate court. By making a decision against a litigant, a trial judge is potentially activating that litigant to mount a strong appeal. In the model I present, the trial judge rationally anticipates an appellate court's review posture, but precisely *how* they respond to appellate review depends on the litigants' behavior.

By emphasizing the role of litigants in the appeal process, this paper fits into a long theoretical literature, largely in the law and economics tradition (e.g., CK 2005, 2006; Daughety and Reinganum 2000; Dewatripont and Tirole 1999; Shavell 1995; Talley 2013). Again, existing results paint an optimistic picture. Two recent "team models" of the judicial hierarchy—CK (2005, 2006)—each demonstrate how litigant-initiated appeals can minimize the number of errors in a three-tiered judicial hierarchy with unbiased judges. The core idea in both articles is that costly appeals serve as signals to judges about the liability of a defendant. By strategically choosing whether to appeal, litigants end up revealing to judges their private information about the case facts.

The analysis in this paper helps clarify the conditions under which the this normatively desirable feature of the appeals process, i.e. error minimization, will occur. For example, a feature of CK (2005, 2006) is that the merits of a case are exogenously learned by the players. This takes two forms: (1) the defendant knows for sure whether he is liable, and (2) courts learn additional information about the case facts with some exogenous probability. In reality, however, litigation often involves a more complex and uncertain informational environment where the litigants and the judge all have some uncertainty about the merits of the case and may have to work hard to figure them out or make sense of them (Posner 2013). Many team models assume judges do not need to make choices about whether to expend resources to do the work of figuring out the case's merits.

The starting point of the model in this paper is that litigation is a protracted and costly process to establish a case's merits (Hornby 2009; Kim et al. 2009; Resnik 1982). The quality of the information available about a case will depend on how much effort the trial judge puts into managing the litigation and how much effort the litigants put into crafting informative appeals. As a result, the model provides a tractable way to capture the empirical reality that judges work hard on cases and litigants work hard to craft persuasive appeals. Beyond providing verisimilitude to the legal process, this framework allows me to endogenize the information collection process that drives the results in similar models, such as CK (2005, 2006). For example, in CK (2005), whether the hierarchy produces zero errors depends on how good courts are at discovering information about cases, which is exogenous. However, if judges have to *choose* how good of a job to do (at a cost), then the fact that defendants reveal their liability through their appeal decisions might actually undermine the judge's incentive to work hard on a case.

The model here also relaxes two assumptions made in previous team models of the appeals process. CK (2005) assumes that judges do not suffer any cost when reversed. In reality, reversals impose at least two costs on judges since they: (1) force a trial judge to use time and resources to reopen and reconsider a case, and (2) are often a source of professional embarrassment. There is ample evidence suggesting that trial judges wish to avoid being reversed (Choi, Gulati, and

4

Posner 2012; Schanzenbach and Tiller 2007). CK (2006) assumes that litigants are symmetric (i.e., have symmetric information). In reality, asymmetries between litigants is an important fact of litigation (see Galanter 1974). For example, of all the federal civil cases filed between January 1 and December 31, 2016, around 25% featured pro se (self-represented) plaintiffs whereas around 3% featured pro se defendants. Looking solely at civil rights cases, these figures are 26% and 2%, respectively.[1] As I show in the analysis, relaxing both of these assumptions can reverse the conclusion that litigant-driven appeals reduce errors.

More broadly, this paper represents a contribution to at least two additional bodies of research. First, the model contributes to research on learning in the judicial hierarchy (e.g., Beim 2017; Beim, Hirsch, and Kastellec 2014; Cameron, Segal, and Songer 2000; Clark and Kastellec 2013; Clark and Carrubba 2012; Carrubba and Clark 2012; Kastellec 2007; Lax 2012). In particular, the model demonstrates that the institutions of the hierarchy can bias law application even when there is no conflict between upper and lower courts over doctrine. The novel feature that drives this result is the role of litigants. Second, the model contributes to existing research on oversight dynamics across political institutions. Some of the closest analogues to the present model come from studies of bureaucracy. For example, researchers have repeatedly demonstrated the importance of bureaucratic effort and expertise for the quality of policy making (e.g., Prendergast 2007; Gailmard and Patty 2007, 2013). Moreover, recent models of bureaucratic incentives, such as Dragu and Polborn (2013), Turner (2016) and Gailmard and Patty (2017), show how review by a supervising body (such as domestic court or an international human rights court) can affect a bureaucrats' effort, and thus outcomes. The model in this paper is most technically similar to the model of endogenous information acquisition in Gailmard and Patty (2017), which explores how agencies respond to notice-and-comment rule-making under judicial oversight.

In the next section, I describe the model. In the subsequent sections, I first demonstrate how

---

[1]These statistics are derived from the Federal Judicial Center's Integrated Database available at `https://www.fjc.gov/research/idb/`.

litigant-driven appeals can generate biased outcomes and more errors. Then, I show how random assignment of cases to judges may exacerbate these problems. Finally, I discuss the implications of the model and conclude.

Model

I analyze a model of a stylized civil case between two litigants, a plaintiff P and a defendant D, initially heard by a trial judge T, and then appealed to an appellate court A. I will sometimes refer to the players using pronouns: "she" for the trial judge, "he" or "them" for the litigants, and "it" for the appellate court.

Cases   I assume that the trial judge and the appellate court agree that there is a correct outcome for every case. Since there is no dispute between them about the relevant law to apply, the trial judge is unbiased (relative to the appellate court). While this assumption is substantively unrealistic, it allows me to more cleanly illustrate the mechanisms that can induce unbiased judges to make biased decisions. Indeed, it serves as a "best case scenario" for assessing the bias-inducing features of judicial institutions. Formally, the case's "merits" are represented by a state of the world $\omega \in \{P, D\}$, indicating whether the law and the facts support the plaintiff or the defendant.

The core problem is that neither court may be certain which party should prevail. Uncertainty over the merits of the case incorporates the idea that a case can be "stronger" for the plaintiff or the defendant even though it may not be completely clear what the outcome should be. As I describe in the game sequence below, litigation serves as a vehicle for the courts to learn more about the merits of the case. The case's merits are a function of both law and facts, so ω can be micro-founded using a case space approach (Lax 2011) where there is uncertainty over facts, law or both. However, doing so adds little insight since the courts are not making new law and do not disagree about existing law.

Sequence   I treat the pre-trial bargaining process between the litigants as a black box, and assume that this bargaining process provides some information about whether the plaintiff or defendant has the better case. Formally, the game begins with a public signal being revealed $t \in \{P, D\}$, which provides (noisy) information about whether the plaintiff or defendant has the stronger case. After receiving this signal, the players all have a belief that $\omega = P$ with probability $0 < \pi < 1$ and $\omega = D$ with probability $1 - \pi$.

It is important to emphasize that the model does not preclude the litigants from engaging in a complicated strategic interaction prior to the trial judge's active involvement in the case. As in other models such as CK (2005), I opt to set these considerations aside since my primary focus in on the incentives that appeals generate for trial judges. Thus, π summarizes how all the players assess the strength of the plaintiff's case at the point where the game begins. I assume the litigants find it worthwhile to continue litigating due to remaining uncertainty about the merits. In the federal courts, the overwhelming majority of cases end without a judgment. For example, between 1996 and 2016, 70% of district court cases ended without a final judgment.[2] A more complicated set up could capture settlements and dismissals, although this is not the primary focus of this paper.

The trial judge plays an important role during the litigation. Specifically, she chooses a level of effort to invest in managing the litigation, $e \in [0, 1]$. Her effort can help improve the quality of information that is learned during litigation. The formal details are as follows. With probability e, additional information is publicly revealed about the merits, $\tau = P$ or $\tau = D$, and with probability $1 - e$ no additional information is revealed about the merits, $\tau = \phi$. Throughout, I use the symbol ϕ to indicate when no additional information is learned. If additional information is revealed ($\tau = P$ or $\tau = D$), it is not perfectly accurate. With probability $\varepsilon < 1/2$, τ is an incorrect signal: $\tau \neq \omega$. With probability $1 - \varepsilon$ it is a correct signal: $\tau = \omega$. Both t and τ are observed by all players. If the trial judge's effort is successful at providing additional useful information, then the players form a

---

[2]These statistics are derived from the Federal Judicial Center's Integrated Database available at `https://www.fjc.gov/research/idb/`. A case was terminated without a final judgment if the `DISP` variable took a value of 0-3 or 10-14.

revised belief about the strength of the plaintiff's case, which I will label π_τ. Formally:

$$\pi_D \equiv \frac{\varepsilon(1-\pi)}{\varepsilon(1-\pi)+(1-\varepsilon)\pi} < \pi \qquad \pi_P \equiv \frac{(1-\varepsilon)\pi}{(1-\varepsilon)\pi+\varepsilon(1-\pi)} > \pi \qquad \pi_\phi = \pi$$

So, after observing $\tau = P$, the players are more convinced that the merits support the plaintiff than they were before, and vice versa after observing $\tau = D$. Although the trial judge's effort may help to improve the quality of the information available when she makes her judgment, there are limits to what she can do. There is already information available (through the signal t) that is provided by the litigants before the trial judge makes any effort. If, for example, the signal t is very strong (π is very high or very low), then the judge's effort won't provide much new insight on the case.

Substantively, the judge's effort can take many forms. A trial judge may hire experts on behalf of the court to evaluate technical information, or she may conduct her own independent research on legislative facts (for a model that incorporates legislative facts, see Beim 2017).[3] Even beyond efforts to supplement the information the litigants provide, the trial judge's efforts may be helpful for eliciting additional information from litigants or for drawing attention to the most informative aspects of the factual record. For example, during litigation in the Northern District of California on the legality of California's ban on same-sex marriage, Judge Vaughn Walker personally oversaw the pre-trial exchange of information, including privately reviewing documents to determine their probative value. After exerting effort on the case (or not), the judge then makes a judgment, $x \in \{D, P\}$, where $x = D$ is a decision in favor of the defendant and $x = P$ is a decision in favor of the plaintiff.[4]

---

[3] The notes of the advisory committee on Rule 201 of the Federal Rules of Evidence provide that a judge may rely on "legislative facts" that are not directly derived from a case's factual record but that are relevant for the resolution of the case.

[4] I model the trial judge's effort in the case as a one-shot process, so that she must make a decision on the case after receiving the signal τ. In reality, a trial judge may make repeated attempts to figure out the merits of a case. As long as there is an increasing marginal cost associated with exerting additional effort (as assumed) and/or delay costs, then including additional attempts to learn the case merits would require additional formal analysis without altering the model's results or providing new insights.

The losing litigant decides how much effort to put into an appeal, $a_L \in [0, 1]$.[5] A successful appeal must convince the appellate court that the lower court made an error, i.e. that $x \neq \omega$. If the trial judge actually did make an error, then exerting more effort to mount a high quality appeal may enable the losing litigant to convince the appellate court of this. I model this idea by assuming that the litigant's effort increases the probability that he discovers a mistake that he can easily show to the appellate court.

I model the litigants' decision calculus differently that other similar models of judicial hierarchy, such as CK (2005, 2006). I assume that the litigants do not have additional private information about the merits once the initial public signal t is revealed. As a result, the losing litigant's decision about whether to appeal is essentially a decision about how many resources to devote to trying to find a correctable error that the appellate court will consider grounds for a reversal. Substantively, this captures the fact that American court procedures favor a relatively open exchange of private information via, for example, liberal discovery rules (Kane 2013). It also captures the notion that a litigant's ability to persuade an appellate court to reverse is tied to the amount of work he does to build his appeal (writing briefs, participating in oral arguments, etc.). Formally, this assumption is not a substantial departure from similar models since the trial judge typically moves first and the pre-trial bargaining is treated as black box. What is important (and a feature of many of these models, including the one here) is that litigants have private information *when they appeal*.

Formally, the losing litigant receives a signal $m_L \in \{\omega, \phi\}$, where ω is learned with probability a_L and ϕ is learned with probability $1 - a_L$. The information learned by the litigant is private, so he has to decide whether to communicate it to the appellate court or conceal it. The litigant does so by writing a brief, which I denote as b_L. If the litigant successfully learns about the merits of the case as a result of his effort ($m_L = \omega$), then he can choose to tell the appellate court what he learned, $b_L = \omega$ or conceal it, $b_L = \phi$. If the litigant receives no information as a result of his effort ($m_L = \phi$), then

---

[5]In this model, it does not matter whether this (or the judge's) effort is observed. In a perfect Bayesian equilibrium, all players must hold correct beliefs about the other players' strategies.

it is not possible to communicate any information so $b_L = \phi$ is the only brief that can be written. I assume that the litigant cannot fabricate information in the brief.[6] Substantively, this could reflect prohibitively high sanctions for perjury combined with a high probability of detection.

Assumption 1 (no fabrication)**.** *Information about the case merits cannot be fabricated.*

Finally, the appellate court decides whether to affirm or reverse the trial court, $r \in \{0, 1\}$.

Payoffs The appellate court gets higher utility when the outcome of the case reflect the merits than when it does not, and I normalize these utilities to 1 and 0, respectively. Given that appellate review is not discretionary for the federal circuit courts, I assume that any cost of review is a sunk cost and thus plays no role in the court's decision-making. The trial judge is unbiased and gets an exogenously fixed benefit $\delta > 0$ from resolving the case in a manner consistent with the case merits. Again, the unbiased judge represents a "best case scenario" for studying institutionally generated bias in the judicial system, since it ensures that the judge has no ex ante incentive to tilt her decision making in favor of either litigant.

The judge also bears two costs. The first is a cost for effort on the case, $\frac{1}{2}c_T e^2$, where $c_T > 0$. This cost reflects both the concrete costs associated with case management and as well as the opportunity cost associated with devoting extra time and energy to the present case and not to other cases. The second is a cost she incurs whenever she is reversed, $k > 0$. This cost reflects the concrete costs associated with reversals (e.g., having to reopen a case) as well as her reputational concerns (e.g., professional embarrassment of being reversed).

The litigants prefer the case to be resolved in their favor and they receive a benefit normalized to one if they do (and zero otherwise). The losing litigant, $L \in \{P, D\}$, pays a cost $\frac{1}{2}c_L a_L^2$ for mounting an appeal, where a_L indicates the amount of effort he devotes to the appeal. I assume

---

[6]Relaxing this assumption to allow the litigant to report D (P) when he receives a signal $m_L = P$ ($m_L = D$) would complicate the analysis. Due to signaling dynamics, it would become harder to sustain an equilibrium where the litigant writes a truthful brief. This would render the *contents* of a litigant's brief relatively unimportant. While this may be reasonable assumption in come contexts, I do not explore it here.

there is no cost associated with filing the appeal, so all litigants appeal. However, they may choose to file a pro forma appeal and put no effort into it (at no cost): $a_L = 0$. I assume the litigant never writes a brief that undermines his own case (i.e., either $b_P = D$ or $b_D = P$).

Assumption 2 (damaging brief)**.** *The appealing litigant never writes a brief that undermines its own case. Formally, $b_P \in \{P, \phi\}$ and $b_D \in \{D, \phi\}$.*

Because the cost parameters c_P and c_D represent both the opportunity costs of the litigants and the resources available to the litigants for building a strong appeal, they also capture, in a sense, the relative power of the litigants. I define the litigant with lower costs as the "more powerful" litigant. To simplify the exposition, I assume that the defendant is the more powerful litigant, $c_P > c_D$. The results do not depend on this assumption, since a symmetric and qualitatively identical set of results emerge if it were reversed.

Finally, I make the following assumption about the litigants' costs, which aids the analysis by ruling out implausible equilibria where the effort of the litigant is so high that they face no serious trade-offs about how to allocate resources.

Assumption 3 (litigant costs)**.** *The litigants' costs are sufficiently high. Formally, $c_P > c_D > 1$.*

Let $\mathbf{1}_y$ be an indicator function taking a value of one when the post-appeal outcome of the case is y. Then the players' utility functions over the outcomes of the game are

$$u_A = \mathbf{1}_\omega \qquad\qquad u_T = \delta\mathbf{1}_\omega - kr - \frac{c_T}{2}e^2 \qquad\qquad u_L = \mathbf{1}_L - \frac{c_L}{2}a_L^2$$

Strategies and Equilibrium   I derive perfect Bayesian equilibria in pure strategies. The trial judge's choice variables are e and x, specifying a level of effort managing the litigation and a judgment. The (losing) litigant's choice variables are a_L and b_L, specifying a level of effort and what kind of brief (informative or not) to submit to the appellate court. The appellate court's choice variable is r, which specifies whether the appellate court affirms or reverses the trial court's judgment. I adopt the following assumptions that restrict the possible equilibria of the game.

Assumption 4 (indifference)**.** *When indifferent, the judge rules in favor of the defendant and the appellate court's reversal strategy favors the defendant.*

Assumption 5 (deference to trial judge)**.** *If there is an equilibrium in which the appellate court defers to any decision by the trial judge, then it does so.*

Assumption 4 rules out uninteresting equilibria driven by players' indifference over outcomes. This assumption is only in operation for knife-edge regions of the exogenous parameter space and is thus unproblematic. Assumption 5 has more bite. It guarantees that, whenever there are multiple equilibria, the appellate court will use a strategy that is "maximally deferential" to the trial judge. I focus on these equilibria because they are closest in spirit to team models, and they reflect the substantive reality of U.S. federal court procedures, which afford wide deference to trial judges' decisions. In the absence of strong institutional rules or norms that induce widespread deference to trial judge decisions, this assumption may be violated. As a result, the analysis in this paper corresponds to situations, as in the U.S., where appellate courts defer to lower courts unless they have very good reason not to.

In the next sections, I describe the logic of the main results. All formal results and proofs are collected in the Online Appendix.

Limits of Litigant-Driven Appellate Review

In this section, I describe the players' optimal strategies by working backward through the model. The analysis will demonstrate the potential limits of litigant-driven appellate review.

The appellate court's optimal reversal strategy is straight forward: it always reverses decisions it believes to be incorrect and affirms decisions it believes to be correct. The appellate court's belief about whether a decision is correct is based on the information it has available to it. Most obviously, if the losing litigant definitively demonstrates that the trial court's judgment is erroneous (by writing a brief of the form $b_L = \omega \neq x$), then the appellate court reverses because it knows

for sure that an error was made. In the model, reversals will sometimes occur because the losing litigant will occasionally finds an error and tells this to the appellate court.

However, if the losing litigant does not definitively demonstrate the trial judge made an error (by writing a brief of the form $b_L = \phi$), then the appellate court's belief is less certain. In this situation, it can do one of four things: reverse any decision, affirm any decision, only reverse decisions for the plaintiff, or only reverse decisions for the defendant. Using Assumption 5, I focus on equilibria where the appellate court affirms any decision by the trial court as long as the appellate court has sufficient uncertainty about which party should prevail.[7]

A losing litigant who discovers an error by the trial judge during its appeal never has an incentive to conceal that information since it leads to a reversal and a final disposition in his favor. In light of this, the losing litigant has to decide whether to expend resources to mount a strong appeal to discover an error. The prospect of a reversal makes it worthwhile for him to exert *some* effort to strengthen his appeal. Formally, the interim expected utility of the losing litigant (who can be either the plaintiff P or the defendant D) is:

$$U_P = a_P \pi_\tau - \frac{1}{2} c_P a_P^2 \qquad\qquad U_D = a_D(1 - \pi_\tau) - \frac{1}{2} c_D a_D^2$$

where π_τ is the litigant's belief that $\omega = P$, conditional on τ. The utility functions above capture each litigant's strategic calculations when deciding how much effort to put into his appeal. First, effort allows him to discover, with probability a_L, whether the trial judge made an error. For example, if the plaintiff lost ($L = P$), then the trial judge made an error with probability π_τ, which the plaintiff discovers with probability a_P. Hence, the plaintiff gets a benefit $a_P \pi_\tau$ for exerting effort a_P. Second, he pays a cost for that effort. For example, if the plaintiff lost, this is $-\frac{1}{2} c_P a_P^2$.

The litigant's interim expected utility function therefore captures the idea that the optimal amount of effort to put into an appeal will balance the potential benefit of discovering an error

---

[7]As I describe in the Online Appendix, this situation arises when $\pi_\tau \in (\underline{\pi}, \overline{\pi}]$, where $0 < \underline{\pi} < 1/2 < \overline{\pi} < 1$.

by the judge with the cost of mounting the appeal. Formally, it can be calculated by maximizing the functions above: $a_P^* = \pi_\tau/c_P$ and $a_D^* = (1 - \pi_\tau)/c_D$.

Judge's Biased Decision Rule

When the trial judge issues her judgment, she is never completely certain whether the merits support the plaintiff or the defendant. This is for two reasons. First, her effort may not be high enough for her to learn more about the case than what she already knew when she began presiding over the case. Second, even if her effort helps her learn more, what she learns can be inaccurate. Specifically, with some probability ε, the additional information she learns about the case is incorrect.

Given that she still has some uncertainty about the merits of the case when she has to make a judgment, she rules in favor of the defendant if her expected utility of doing so is greater than or equal to her expected utility of ruling in favor of the plaintiff. Moreover, since she anticipates the litigant's appeal, these expected utilities depend on the losing litigant's strategy. Formally, trial judge rules for the defendant if and only if

$$\underbrace{(1 - \pi_\tau)\delta + \pi_\tau(\delta - k)a_P}_{\text{rules for defendant (against plaintiff)}} \geq \underbrace{\pi_\tau\delta + (1 - \pi_\tau)(\delta - k)a_D}_{\text{rules for plaintiff (against defendant)}}$$

where π_τ is the trial judge's belief that $\omega = P$. Given the information structure of this game, her belief takes one of three values: $\pi_\tau = \pi$ (she didn't learn anything after her effort), $\pi_\tau = \pi_P$ (she learned $\tau = P$), and $\pi_\tau = \pi_D$ (she learned $\tau = D$). Conversely, the judge rules in favor of the plaintiff when the inequality is strictly satisfied in the other direction. Substituting the equilibrium values for a_D^* and a_P^* from above and rearranging, the judge rules for the defendant if and only if:

$$(1 - 2\pi_\tau)\delta + (\delta - k)\left(\frac{\pi_\tau^2}{c_P} - \frac{(1 - \pi_\tau)^2}{c_D}\right) \geq 0 \tag{1}$$

The left hand side is strictly positive at $\pi_\tau = 0$, strictly negative at $\pi_\tau = 1$ and strictly decreasing

in π_τ. Therefore, there is a unique threshold $\tilde{\pi}$ such that (1) holds with equality. Intuitively, this threshold represents how certain the judge has to be that the merits favor the plaintiff before being willing to rule in favor of the plaintiff.[8] That is, it's her **decision rule**: she rules in favor of the defendant if $\pi_\tau \leq \tilde{\pi}$ and in favor of the plaintiff if $\pi_\tau > \tilde{\pi}$.

A trial judge who uses an **impartial decision rule** would base her judgment solely on what she believes about the merits, and not on other considerations. Formally, she would rule in favor of the defendant if she believes $\pi_\tau \leq 1/2$ and in favor of the plaintiff if she believes $\pi_\tau > 1/2$. In the model, however, the trial judge bases her decision on the threshold $\tilde{\pi}$, which is not generally equal to $1/2$. Therefore, in equilibrium, the trial judge ends up biasing her decisions to favor one of the litigants. This is an institutionally generated bias since the trial judge does not personally favor either of the litigants. Moreover, it is the fact that appeals are *litigant-driven* that generates the bias, not appellate review per se. Without strategic litigants, her decision rule would be impartial: $\tilde{\pi} = 1/2$. Whether this institutional bias favors the plaintiff or the defendant overall depends on how reversal averse the trial judge is.

Proposition 1. *The trial judge's decision rule has the following properties:*

- *If $k < \delta$, then her decision rule is biased in favor of the less powerful litigant, $\tilde{\pi} < 1/2$.*

- *If $k > \delta$, then her decision rule is biased in favor of the more powerful litigant, $\tilde{\pi} > 1/2$.*

- *If $k = \delta$, then her decision rule is impartial, $\tilde{\pi} = 1/2$.*

Moreover, the bias in her decision rule becomes weakly larger as $|\delta - k|$ increases.

Figure 1 depicts this result by plotting the trial judge's decision rule, represented by the threshold $\tilde{\pi}$, as a function of her reversal aversion, k. If $k < \delta$, then the institutional bias favors the weaker litigant (here, the plaintiff, depicted in region I). In this situation, the trial judge cares very little about being reversed. She rules against the powerful party in the absence of information because

---

[8]The proof of Lemma A4 in the Online Appendix explicitly characterizes $\tilde{\pi}$ as a function of the parameters in condition (1).

she knows the powerful party will mount a more effective appeal. After all, she values the losing party's effort to figure out the merits. If $k > \delta$, then this institutional bias favors the more powerful litigant (here, the defendant, depicted in region II). In this situation, the trial judge dislikes reversals sufficiently that she minimizes her chance of being reversed by stacking the deck in favor of more powerful litigants, who would otherwise mount stronger appeals.

Figure 1: *The judge's decision rule is not generally impartial, and as her reversal aversion increases, the decision rule she uses increasingly favors the more powerful litigant*

(assuming: $\delta = 1$, $c_T = 1.1$, $c_D = 1.1$, $c_P = 2$)
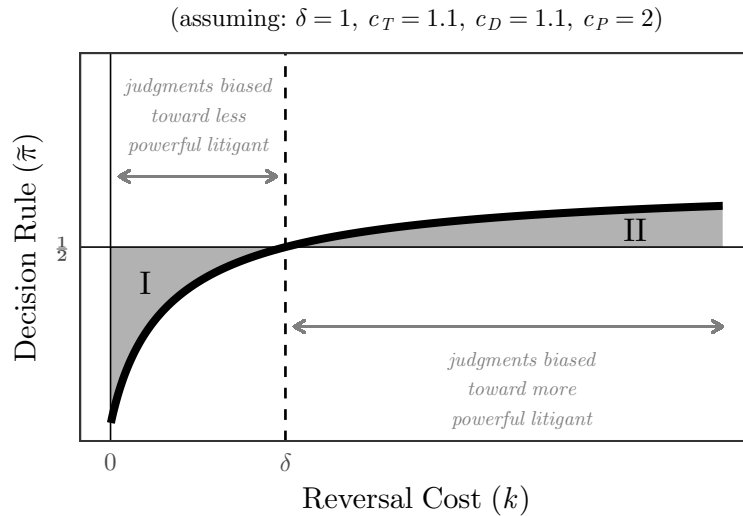
The fact that the trial judge uses a biased decision rule does not mean that she flagrantly ignores the law when she makes a judgment. Recall that she has some uncertainty about the case's merits, as represented by π_τ (her belief that $\omega = P$). She biases her decisions when she is the *most uncertain* about whether the plaintiff or defendant should prevail. Formally, she does so whenever π_τ is relatively close to $1/2$, the point where she thinks there is a 50-50 chance that the plaintiff's case is stronger. Again, this is apparent in Figure 1 since the regions that yield biased outcomes (I and II) are close to the $1/2$ line.

In the analysis above, the trial judge's optimal decision rule relies on her presumption that the appellate court will affirm her judgment in the absence of an informative appeal from the litigant.

16

However, the appellate court can observe that the trial judge is using a biased decision rule. Surprisingly, I show in Lemma A4 in the Online Appendix that as long as her decision rule is not *too* biased, then there is an equilibrium in which the appellate court affirms judgments even when it *knows* the trial judge is using a biased decision rule. This is because an uninformative appeal sends a mixed signal to the appellate court. On the one hand, the litigant may have had a hard time finding a "smoking gun" that demonstrates an error was made. On the other hand, he could be concealing some unfavorable information he discovered that would undermine his case. Since the litigant is motivated to win, the appellate court is skeptical of him. This skepticism gives the trial judge some room to maneuver without provoking a reversal. So, not only can litigant-driven appeals induce the trial judge to bias her judgments, they can also give her the flexibility to do so while avoiding reversal.

Before proceeding, compare this result with CK (2005), which demonstrates that lower court judges will systematically rule against litigants with more information so that a higher court may potentially elicit this information. In the informational environment here, this is isomorphic to ruling against the powerful litigant, who is better able to detect errors. Proposition 1 also contains this result, but it *only* occurs for the case where $k < \delta$. The fact that CK (2005) assumes that $k = 0$ thus turns out to be important for the conclusion that hierarchy encourages error correction. I show here that if a judge is sufficiently reversal averse, then she may do exactly the opposite: rule against the weaker litigant in order to make it *more difficult* for errors to be detected.

Judge's Effort Minimization

Next, I consider the judge's decision about how much effort to invest in managing the case. For brevity, I omit details about the derivation, but interested readers should refer to Lemma A6 in the Online Appendix. Intuitively, she faces a trade-off: effort helps her make a more accurate decision (which she likes), but it is costly. Let $x(\phi)$ be the judgment the trial judge would make if her effort is not successful at improving the quality of information available, i.e. if $\tau = \phi$. Then, given her

17

choice of $x(\phi)$, if the trial judge finds it optimal to exert some amount of effort $e \in (0, 1)$, then it is as follows:[9]

$$
e^* = \begin{cases} \dfrac{1}{c_T}\left[(1 - \pi - \varepsilon)\delta + (k - \delta)\left(\dfrac{(1 - \pi)^2}{c_D}(1 - \varepsilon) - \dfrac{\pi^2}{c_P}\varepsilon\right)\right] & \text{if } x(\phi) = P \\[3ex] \dfrac{1}{c_T}\left[(\pi - \varepsilon)\delta + (k - \delta)\left(\dfrac{\pi^2}{c_P}(1 - \varepsilon) - \dfrac{(1 - \pi)^2}{c_D}\varepsilon\right)\right] & \text{if } x(\phi) = D \end{cases}
$$

Suppose that ε is sufficiently low, so that the trial judge's effort yields sufficiently accurate information about the case. Then, litigant-driven appellate review affects this effort in two important ways. First, the possibility that the litigant will detect a reversible error changes the amount of effort that she finds optimal. If $k > \delta$, the judge dislikes reversals so much that she may exert more effort to prevent making an error. To see this, notice that when $k > \delta$ (and ε is small), e^* increases as k increases. In this situation, litigant-driven appeals provide a standard "agency benefit" since it provides appellate courts with a way to mitigate a trial judge's incentive to shirk. However, if $k < \delta$ (and ε is small), then this conclusion no longer holds. In that situation, then appellate review actually *reduces* her effort since she now counts on the losing litigant to provide additional information if she makes an error.

Proposition 2. *If ε is sufficiently low, then litigant-driven appellate review has the following effect on the trial judge's equilibrium effort:*

- *If $k < \delta$, then her effort is strictly lower than without litigant-driven appellate review.*

- *If $k > \delta$, then her effort is strictly higher than without litigant-driven appellate review.*

- *If $k = \delta$, then litigant-driven appellate review has no effect on the trial judge's equilibrium effort.*

The second way that litigant-driven appeals affects effort is more subtle. Recall that the previous analysis establishes the trial judge uses a decision rule that generates biased outcomes. Relative to

---

[9]Note: as Lemma A6 in the Online Appendix shows, it is possible for the trial judge to choose $e = 0$ or $e = 1$. The fact that these corner solutions are possible does not provide any interesting insights.

an impartial decision rule, this allows her to exert weakly lower effort. When $k > \delta$, she biases her decision-making to favor the powerful litigant and thus worries less about the losing litigant discovering an error. When $k < \delta$, she biases her decision-making against the more powerful litigant and thus is more confident that an error will get discovered by the litigant. Either way, this allows her to lower her effort level relative to what she would have done if she used an impartial decision rule.

Proposition 3. *If $\delta \neq k$, then the trial judge's equilibrium effort is weakly lower than if she used an impartial decision rule. Moreover, it is strictly lower for all $\pi \in (\max\{\frac{1}{2}, \tilde{\pi}\}, \min\{\frac{1}{2}, \tilde{\pi}\}]$.*

I depict this result graphically in Figure 2, which plots e^* as a function of the strength of the plaintiff's case (π) for some fixed parameter values. The thick line is the equilibrium level of effort.

Figure 2: *Equilibrium effort depends on π and is weakly lower than effort under an impartial decision rule (which is not optimal)*

(assuming: $\delta = 1$, $k = 6$, $c_T = 1.1$, $c_D = 1.1$, $c_P = 2$, $\varepsilon = 0.2$)
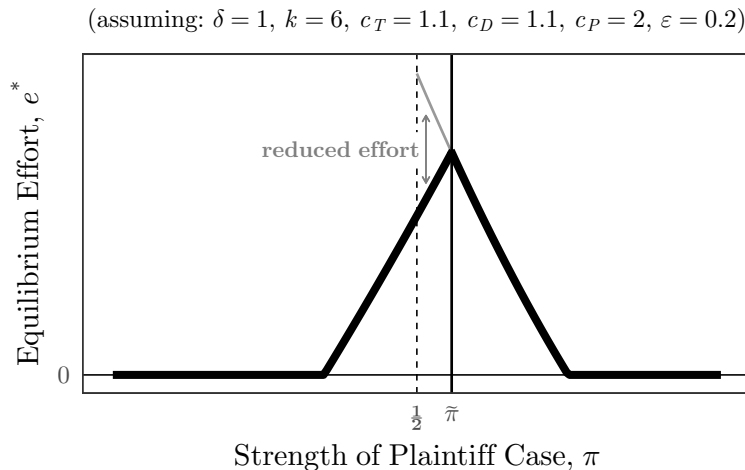
From the analysis above, her judgment shifts from pro-defendant to pro-plaintiff at $\tilde{\pi}$, which in Figure 2 is biased in favor of the more powerful litigant. The fact that she can shift her judgment at $\tilde{\pi}$ means that she minimizes the level of effort that she exerts. Of particular interest is the region $1/2 < \pi < \tilde{\pi}$ where she makes biased decisions. If she made decisions according to an impartial

decision rule instead, her effort in this interval would be the thin gray line. Because she biases her decisions, her effort ends up on the thick black line, which is strictly lower.

Again contrast this result with CK (2005). In that model, the judges' ability to learn about a case is fixed and exogenous. The results in this paper suggest that when $\delta > k$, trial judges actually have an incentive to *lower* the probability they learn about the case merits (i.e. reduce e) if learning requires costly effort and if they know they can rely on the losing litigant to find their errors.

Case Assignment and Diversity on the Bench

The analysis has so far established two important findings: (1) an institutional bias in a trial judge's decision-making may emerge even when the judge is not biased, and (2) trial judges minimize their effort through their strategic choice of who wins a case. These two findings have important implications for the way that the aggregate set of cases gets assigned to, and thus resolved by, a pool of heterogeneous judges. In this section, I explore this issue formally. To focus on the key insights, I assume that $\varepsilon = 0$ so that anything the trial judge is able to learn about the case is accurate (i.e., $\tau = P$ if and only if $\omega = P$).[10] As before, she still may not learn anything at all ($\tau = \phi$).

If trial judges bias their decision making and minimize their effort as described in the previous section, then this suggests the judicial hierarchy may not always minimize the number of errors made. Formally, in the equilibrium described in the previous section, the probability that a case will be resolved without error can be calculated by the following, which I label $\xi(\pi, \delta, c_T)$:[11]

$$\xi(\pi, \delta, c_T) = \begin{cases} e^* + (1 - e^*)(1 - \pi(1 - a_P^*)) & \text{if } x(\phi) = D \\ e^* + (1 - e^*)(1 - (1 - \pi)(1 - a_D^*)) & \text{if } x(\phi) = P \end{cases}$$

---

[10]The results in this section hold as long as ε is sufficiently low. However, since the exposition becomes much more cumbersome, I focus on the simplest case where $\varepsilon = 0$.

[11]Recall, the trial judge and appellate court do not disagree about which law to apply, so in this model there is a clear notion of what constitutes a "correct" decision.

I now use this quantity to explore the aggregate ramifications of the model when a set of cases enters the trial court and there is a set of judges available to hear those cases.

Suppose that there is a pool of trial judges, each of whom could be assigned to a case. Each judge differs on two dimensions: the degree to which they are intrinsically motivated by different kinds of cases and their cost for exerting effort. Formally, there is heterogeneity among judges in their δ and c_T parameters. I make the substantively reasonable assumption that a judge's δ and c_T are negatively correlated. That is, judges who are more interested in a specific kind of case (higher δ) are also more effective at managing those kinds of cases (lower c_T). Substantively, this captures the idea that judges will tend to develop more expertise on issues that interest them than on issues that do not interest them.[12] Then, in a judicial system with K issues that can come before the courts, each judge i can be characterized by a vector, $((\delta_1^i, c_1^i), (\delta_2^i, c_2^i), ..., (\delta_K^i, c_K^i))$, indicating their degree of intrinsic motivation and cost of managing cases for each issue, $k \in \{1, 2, ...K\}$. (Note: i superscripts index judges and k subscripts index issue areas.)

To focus on the core substantive lessons, I examine the simplest situation where there are two kinds of judges $i \in \{1, 2\}$, two issues $k \in \{1, 2\}$, two levels of issue motivation $\delta_k^i \in \{0, \bar{\delta}\}$ (where $\bar{\delta} > 0$), and two levels of effort cost, $c_k^i \in \{c_L, c_H\}$ (where $0 < c_L < c_H$). Moreover, I assume that the two kinds judges are interested in different issues:[13]

$$\underbrace{((\delta_1^1 = \bar{\delta}, c_1^1 = c_L), (\delta_2^1 = 0, c_2^1 = c_H))}_{\text{issue 1 judges}} \qquad \underbrace{((\delta_1^2 = 0, c_1^2 = c_H), (\delta_2^2 = \bar{\delta}, c_2^2 = c_L))}_{\text{issue 2 judges}}$$

To emphasize which issues each kind of judge cares more about, I refer to the judges with preferences on the left "issue 1 judges" and judges with preferences on the right "issue 2 judges." For example, issue 1 judges may be very interested in civil rights cases but not in securities litiga-

---

[12]In light of the equilibrium dynamics described in the previous section, this notion should be uncontroversial. The formal analysis above demonstrates that if $\varepsilon = 0$, judges exert more effort as δ increases. Then, judges learn more about cases where they have higher δs, potentially making it easier to learn about similar cases in the future. While a full analysis of the endogenous process by which c_T changes is outside the scope of this paper, this logic provides a rationale for assuming a negative correlation between δ and c_T.

[13]The insights would emerge from a more complicated setup, although with loss of clarity.

tion, while issue 2 judges may be very interested in securities litigation, but not civil rights cases. Each kind of judge will make more or fewer errors depending on what kind of case they hear. The probability of an error-free outcome when the assigned trial judge is interested in the case is $\xi(\pi, \bar{\delta}, c_L) \equiv \bar{\xi}$. Alternatively, the probability of an error-free outcome when the assigned trial judge is not interested in the case is $\xi(\pi, 0, c_H) \equiv \underline{\xi}$. The fact that the judges differ with respect to their issue interests implies the following result, which is useful in the subsequent analysis.

Lemma 1. *There are more errors when issue 2 judges hear issue 1 cases than when issue 1 judges hear issue 1 cases, and vice versa. Formally, $\bar{\xi} \geq \underline{\xi}$, where the condition holds strictly if $c_H > \bar{c}_T$ (with $\bar{c}_T$ defined in the proof).*

The result in Lemma 1 comes from two sources, one direct and one indirect. First, the effort of a judge increases as δ increases. Somewhat intuitively: when judges care more about their cases, they work harder on them. Second, when judges care less about cases (e.g., issue 1 judges on issue 2 cases), they bias their decisions toward the powerful litigant (from Proposition 1). This reduces the probability an error will get detected and thus lowers the accuracy of outcomes.

I now describe two case assignment procedures, which are at the extremes of a broader set of potential assignment procedures. Suppose the mix of judges given by q, which is the probability that a given judge is an issue 1 judge. Accordingly, $1-q$ is the probability that a given judge is an issue 2 judge. Moreover, suppose an individual case is an issue 1 case with probability p and an issue 2 case with probability $1-p$. First, suppose that each case is randomly assigned to a judge and that random assignment gives equal weight to each judge. Then, the expected proportion of correct decisions under random assignment, which I label R, is given by $R \equiv p\big(q\bar{\xi}+(1-q)\underline{\xi}\big)+(1-p)\big(q\underline{\xi}+(1-q)\bar{\xi}\big)$.

Second, at the opposite extreme, suppose instead that when cases come into the court system, trial judges volunteer to take them. It is straight forward to observe that, all else equal, an issue 1 judge will choose to take issue 1 cases over issue 2 cases and an issue 2 judge will choose to take issue 2 cases over issue 1 cases. I assume that the volunteer assignment process satisfies two assumptions. First, cases end up on the dockets of judges who have an interest in them as long as

there is room on those judges' dockets. Second, all judges' dockets must be equally populated by cases. Suffice to say, the aggregate effect of this assignment system is that cases will be assigned to judges based on their interest so long as those judges are not overburdened. Before proceeding, note that the analysis below would be unchanged with a slightly more nuanced procedure where judges *specialize* in an issue and have cases randomly assigned to them within their specialization.

Given that there may be substantial differences in the number of issue 1 (or 2) judges than there are issue 1 (or 2) cases and I assume dockets must be equal across judges, a quantity of interest will be $q - p$. If $q - p > 0$, then there are more issue 1 judges than there are issue 1 cases, and vice versa if $q - p < 0$. The knife-edge case of $q = p$ represents a situation where there are just enough issue 1 judges to take all issue 1 cases and just enough issue 2 judges to take all issue 2 cases. Let $m \equiv \max\{q - p, p - q\}$ represent the divergence between the mix of judges and the mix of cases. Substantively, m is the proportion of cases that have to be resolved by "mismatched" judges. Then, the expected proportion of correct decisions under volunteer assignment, which I label V, is given by $V \equiv (1 - m)\bar{\xi} + m\underline{\xi}$.

Having characterized the accuracy of case outcomes under these two assignment procedures, the following result demonstrates an important downside to random assignment.

Proposition 4. *Random assignment of judges to cases leads to fewer accurate decisions than voluntary assignment. Formally, $R \leq V$, and $R < V$ if $c_H > \bar{c}_T$.*

Proposition 4 compares two extreme assignment procedures. More generally, an assignment procedure that allows judges more flexibility to pick their cases will generate fewer errors than random assignment. There are two obvious institutional implications. First, this finding underscores a benefit of specialized courts, since they allow judges to select into a set cases that interest them. What drives Proposition 4 in the fact that a judge's effort is affected by their level of issue interest. Proposition 4 suggests that one way to improve the accuracy of case outcomes is to allow judges to choose cases that interest them. This provides a rationale for why specialized courts are often

focused on certain areas of law that are complex and require a lot of active judicial involvement (e.g., bankruptcy, patents, etc.).

A second institutional implication of Proposition 4 involves judicial selection. In light of the predictions of this model, one optimal way to select judges is to appoint those who are highly motivated on every single issue that could come before the courts. Realistically, though, this is not feasible. Judicial appointments almost always require choices between candidates with differing priorities. One potential response is to appoint judges from a wide variety of backgrounds, and thus with a variety of priorities. Indeed, in recent years, the importance of "diversity on the bench" has been front and center in debates about the composition of the judiciary. One way to represent "diversity on the bench" in the model is the parameter q, the mix of judges on the bench. Under random assignment, a court minimizes the number of errors by having an entirely homogeneous bench since R is monotonically increasing or decreasing in q. Under voluntary assignment, this is is not the case. Since V increases as the mix of judges q approaches the mix of cases p, diversity will reduce errors under voluntary assignment. Moreover, the benefit of voluntary assignment over random assignment (see Proposition 4) actually *increases* as the mix of judges reflects the mix of cases. Figure 3 depicts this graphically.

One interesting issue that I do not explore here is the effect of these dynamics on the mix of cases coming into the courts. Essentially, I assume p is exogenous and unchanged as q changes. An interesting avenue for future research would be to analyze the impact of increasing diversity (i.e., changing q) on litigants' decisions to file cases in court.

The result in Proposition 4 applies to a model of adjudication as described in the analysis of the previous sections. Random assignment, however, also speaks to a set of concerns about judicial bias, which is not captured in the core model. Specifically, if judges volunteer to take cases, they may take cases in which they are biased in favor of one party or another. Random assignment is supposed to prevent this. Importantly, Proposition 4 does *not* say that random assignment is *generally* worse than volunteer assignment. Indeed, the result holds for the model of adjudication analyzed

Figure 3: *Accuracy of adjudication under different case assignment procedures*

(assuming: $p = 0.4$, $\bar{\xi} = 0.5$, $\underline{\xi} = 0.25$)
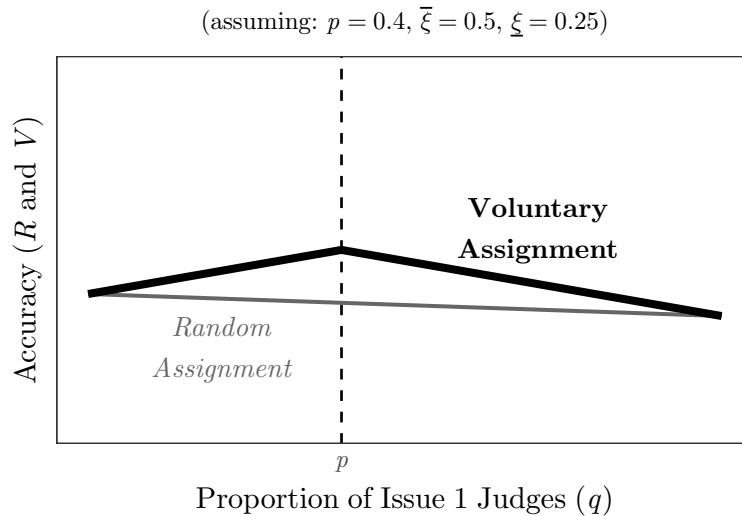
in the previous section. So, to the extent that the dynamics of that model operate in real-world decision-making and judges have heterogeneous preferences over which cases they are interested in, then Proposition 4 establishes an often overlooked disadvantage of random assignment: it puts too many disinterested judges on cases, reducing the overall accuracy of case outcomes since those disinterested judges make more errors.

Discussion

The extent to which the legal system produces biased outcomes is a ripe area of concern among public officials, legal practitioners, academics and outside observers. Indeed, judges routinely insist that personal biases do not influence their decision making. And yet, a complex web of institutions has arisen to insulate the judicial system against the idiosyncratic views of individual judges. My analysis demonstrates that there may be unrecognized downsides to these institutions. In particular, litigant-driven appeals can (1) incentivize and enable a trial judge to make biased decisions and (2) allow a trial judge to reduce her effort in managing a case. Moreover, the random assignment of

judges to cases draws out these negative effects, which could be mitigated with other, more flexible assignment procedures.

In the model, these negative effects emerge despite the fact that I assume trial judges are completely unbiased. This suggests that judicial bias is not a necessary condition for biased outcomes. It also suggests that institutions do not always have the effects that are intended. While I do not explore how these results apply in a context where judges are personally biased, the results suggest that litigant-driven appeals can provide all judges (biased or not) with some leeway to bias their decisions away from established doctrine. In the analysis, the core reason for this is the appellate court's skepticism of losing litigants who are motivated to win cases and who may therefore have an incentive to conceal unfavorable information. This skepticism can exist whether or not a trial judge is biased.

The centrality of litigant behavior provides important lessons for scholars studying the way that the judicial hierarchy affects judges' decision making. Perhaps most directly, the results demonstrate that litigants' behavior is crucial for understanding the full set of incentives that judges face. This is particularly true of trial judges, whose interactions with litigants are frequent, regular and often determinative. Models of judicial hierarchy that speak to the law application purpose of lower courts (such as those involving judges learning about fact patterns and applying law) should explore the way that litigants affect judges' behavior. For example, an interesting avenue for future theoretical work would be to explicitly model the pre-trial bargaining process that is black boxed in this paper (and others, such as CK 2005, 2006). While ample theoretical research has examined this topic (most famously, Priest and Klein 1984), much of it treats judges as either relatively passive or unbiased. As I demonstrate here, the interaction between judge and litigant incentives is not always obvious.

The analysis also suggests that empirical work on the judicial decision-making should be careful to account for the ways that litigants (and especially imbalance between litigants) may affect the range of options available to a judge. While the empirical implications of the model are nuanced, a

couple things stand out. First, the model suggests that a judge's decision rule will depend on *both* the strength of the plaintiff's case (represented by π and tied to case characteristics) *and* resource imbalances between plaintiffs and defendants. While scholars have long recognized the importance of case characteristics in the study of legal outcomes (for example, see Kastellec 2010), less explicit attention has been paid to resource imbalances in empirical studies (but see Songer, Sheehan, and Haire 1999). Second, the model also makes clear that personal (ideological or other) bias is not the only explanation for empirical results showing judges making decisions that systematically favor one kind of litigant. More subtly, the results also suggest an observationally equivalent explanation for empirical findings that certain types of judges rule differently than other types of judges (for example, Boyd 2016).

Why are the results more normatively troubling than previous team models of the judicial hierarchy? Consider CK (2005), which shows how judges' decision rules maximize the chance that errors are detected and reversed on appeal. My analysis suggests those decision rules may be sensitive to the assumption that judges do not care if they are reversed. Relaxing this assumption can completely reverse the decision rule used by a trial court: judges with sufficient reversal aversion actually *minimize* the chance that errors are detected and reversed on appeal. As a result, the normatively desirable conclusion that judges use decision rules that maximize the detection of errors only emerges as a special case of the model in this paper, i.e. when $k < \delta$. Even for this special case, however, the results here do not necessarily generate maximal detection of errors. Unlike CK (2005), I assume that the judge's ability to learn about the case is endogenous and tied to how much (costly) effort she puts into managing the case. Since the trial judge exploits the powerful litigant when $k < \delta$, she ends up *lowering* her effort since she knows the litigant will pick up the slack on appeal. This suggests that this potential "cost sharing" behavior by the trial judge is not really cost sharing: she is shirking by substituting the litigant's effort for her own.

Finally, this paper provides an important exception to the view that team models of the judicial hierarchy show "how the appellate system can mitigate the problem of incorrect decisions by lower

courts" (Kastellec 2017). The analysis demonstrates that sufficient reversal aversion can turn this view on its head. In doing so, it implies that judges with reversal aversion (as here, and as in CK 2006) may not always be on the same team as the appellate judges who oversee them. In those situations, and even when judges are unbiased, there is reason to be skeptical that the appeals process and random assignment encourage error-minimizing, unbiased decision-making. Paradoxically, institutions designed to limit the influence of individual judges can actually magnify their negative influence.

References

Administrative Office of the United States Courts. 2018. *Federal Judicial Caseload Statistics*. https://www.uscourts.gov/statistics-reports/analysis-reports/federal-judicial-caseload-statistics (accessed August 11, 2019).

Beim, Deborah. 2017. "Learning in the Judicial Hierarchy." *Journal of Politics* 79 (2): 591–604.

Beim, Deborah, Alexander V. Hirsch, and Jonathan P. Kastellec. 2014. "Whistleblowing and Compliance in the Judicial Hierarchy." *American Journal of Political Science* 58 (4): 904–918.

Boyd, Christina L. 2015. "The Hierarchical Influence of Courts of Appeals on District Courts." *Journal of Legal Studies* 44 (1): 113–141.

Boyd, Christina L. 2016. "Representation on the Courts? The Effects of Trial Judges' Sex and Race." *Political Research Quarterly* 69 (4): 788–799.

Boyd, Christina L., and James F. Spriggs II. 2009. "An Examination of Strategic Anticipation of Appellate Court Preferences by Federal District Court Judges." *Washington University Journal of Law and Policy* 29:37–81.

Cameron, Charles M., and Lewis A. Kornhauser. 2005. "Decision Rules in a Judicial Hierarchy." *Journal of Institutional and Theoretical Economics* 161 (2): 264–292.

Cameron, Charles M., and Lewis A. Kornhauser. 2006. "Appeals Mechanisms, Litigant Selection and the Structure of Judicial Hierarchies." In James R. Rogers, Roy B. Flemming and Jon R. Bond, eds., *Institutional Games and the U.S. Supreme Court.* Charlottesville, VA: University of Virginia Press, 173–204.

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." *American Political Science Review* 94 (1): 101–116.

Carrubba, Clifford J., and Tom S. Clark. 2012. "Rule Creation in a Political Hierarchy." *American Political Science Review* 106 (3): 622–643.

Choi, Stephen J., Mitu Gulati, and Eric A. Posner. 2012. "What Do Federal District Judges Want? An Analysis of Publications, Citations, and Reversals." *Journal of Law, Economics, and Organization* 28 (3): 518–549.

Clark, Tom S., and Clifford J. Carrubba. 2012. "A Theory of Opinion Writing in a Political Hierarchy." *Journal of Politics* 74 (2): 584–603.

Clark, Tom S., and Jonathan P. Kastellec. 2013. "The Supreme Court and Percolation in the Lower Courts: An Optimal Stopping Model." *Journal of Politics* 75 (1): 150–168.

Daughety, Andrew F., and Jennifer F. Reinganum. 2000. "On the Economics of Trials: Adversarial Process, Evidence, and Equilibrium Bias." *Journal of Law, Economics, and Organization* 16 (2): 365–394.

Dewatripont, Mathias, and Jean Tirole. 1999. "Advocates." *Journal of Political Economy* 107 (1):

1–39.

Dragu, Tiberiu, and Mattias K. Polborn. 2013. "The Administrative Foundation of the Rule of Law." *Journal of Politics* 75 (4): 1038–1050.

Edwards, Harry T. 1998. "Collegiality and Decision Making on the D.C. Circuit." *Virginia Law Review* 84 (7): 1335–1370.

Epstein, Lee, William M. Landes, and Richard A. Posner. 2013. *The Behavior of Federal Judges: A Theoretical and Empirical Study of Rational Choice*. Cambridge, MA: Harvard University Press.

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

Gailmard, Sean, and John W. Patty. 2013. *Learning While Governing: Expertise and Accountability in the Executive Branch*. Chicago: University of Chicago Press.

Gailmard, Sean, and John W. Patty. 2017. "Participation, Process, & Policy: The Informational Value of Politicized Judicial Review." *Journal of Public Policy* 37 (3): 233–260.

Galanter, Marc. 1974. "Why The "Haves" Come Out Ahead: Speculations on the Limits of Legal Change." *Law and Society Review* 9 (1): 95–160.

Hornby, D. Brock. 2009. "The Business of the U.S. District Courts." In David M. O'Brien, ed., *Judges on Judging: Views from the Bench*. 3rd ed. Washington, DC: CQ Press, 88–96.

Kane, Mary Kay. 2013. *Civil Procedure in a Nutshell*. 6th ed. St. Paul, MN: Thomson-West.

Kastellec, Jonathan P. 2007. "Panel Composition and Judicial Compliance on the US Courts of Appeals." *Journal of Law, Economics, and Organization* 23 (2): 421–441.

Kastellec, Jonathan P. 2010. "The Statistical Analysis of Judicial Decisions and Legal Rules with Classification Trees." *Journal of Empirical Legal Studies* 7:202–230.

Kastellec, Jonathan P. 2017. "The Judicial Hierarchy." In *Oxford Research Encyclopedia of Politics*. Oxford University Press.

Kim, Pauline T., Margo Schlanger, Christina L. Boyd, and Andrew D. Martin. 2009. "How Should We Study District Judge Decision-Making?" *Journal of Law and Policy* 29 (83): 83–112.

Lax, Jeffrey R. 2011. "The New Judicial Politics of Legal Doctrine." *Annual Review of Political Science* 14:131–157.

Lax, Jeffrey R. 2012. "Political Constraints on Legal Doctrine: How Hierarchy Shapes the Law." *Journal of Politics* 74 (3): 765–781.

Posner, Richard A. 2013. *Reflections on Judging*. Cambridge, MA: Harvard University Press.

Prendergast, Canice. 2007. "The Motivation and Bias of Bureaucrats." *American Economic Review* 97 (1): 180–196.

Priest, George L., and Benjamin Klein. 1984. "The Selection of Disputes for Litigation." *Journal of Legal Studies* 13:1–55.

Randazzo, Kirk A. 2008. "Strategic Anticipation and the Hierarchy of Justice in U.S. District Courts." *American Politics Research* 36 (5): 669–693.

Resnik, Judith. 1982. "Managerial Judges." *Harvard Law Review* 96:374–448.

Revesz, Richard L. 1997. "Environmental Regulation, Ideology, and the D.C. Circuit." *Virginia Law Review* 83 (8): 1717–1772.

Schanzenbach, Max M., and Emerson H. Tiller. 2007. "Strategic Judging Under the U.S. Sentencing Guidelines: Positive Political Theory and Evidence." *Journal of Law, Economics, and Organization* 23 (1): 24–56.

Shavell, Steven. 1995. "The Appeals Process as a Means of Error Correction." *Journal of Legal Studies* 24 (2): 379–426.

Sherman, Mark. 2018. "Roberts, Trump Spar in Extraordinary Scrap over Judges." *AP*, November 21, 2018. https://apnews.com/c4b34f9639e141069c08cf1e3deb6b84 (accessed August 11, 2019).

Songer, Donald R., Reginald S. Sheehan, and Susan Brodie Haire. 1999. "Do the "Haves" Come out Ahead over Time? Applying Galanter's Framework to Decisions of the U.S. Courts of Appeals, 1925-1988." *Law and Society Review* 33 (4): 811–832.

Talley, Eric. 2013. "Law, Economics, and the Burden(s) of Proof." In Jennifer H. Arlen, ed., *Research Handbook on the Economic Analysis of Tort Law*. Northampton, MA: Edward Elgar, 305–329.

Turner, Ian R. 2016. "Working Smart and Hard? Agency Effort, Judicial Review, and Policy Precision." *Journal of Theoretical Politics* 29 (1): 69–96.