Crafting the Law:
How Opinion Content Influences Legal Development*

Michael J. Nelson
Assistant Professor
Department of Political Science
The Pennsylvania State University
mjn15@psu.edu

Rachael K. Hinkle†
Assistant Professor
Department of Political Science
University at Buffalo, SUNY
rkhinkle@buffalo.edu

September 14, 2016

Why are some judicial opinions widely discussed while others languish in obscurity? We theorize that opinions that can be understood efficiently are discussed, expanded, and contracted more frequently. Additionally, more persuasive precedents will tend to be discussed and expanded more frequently, while less persuasive precedents will be narrowed more often. These effects should also depend upon institutional context. We compile an original dataset of 7,604 search and seizure opinions written by federal circuit and state high courts and all 676,011 judicial citations to those cases. Using these data, we track vertical influence up and down the judicial hierarchy and evaluate horizontal influence both within the precedent's jurisdiction and across jurisdictional lines. There is evidence that efficiency, persuasiveness, and institutional context all play a role in how much an opinion impacts legal development.

---

*Data and code necessary to replicate the analyses in this paper will be available on the authors' websites upon publication.

†We thank Doug Rice and Tom Clark for comments on an earlier version of this manuscript, presented at the 2015 meeting of the Midwest Political Science Association and the 10th Annual Conference on Empirical Legal Studies.

Judicial opinions have both short-term and long-term effects. In the short-term, judicial decisions resolve a controversy between the parties to the case, adjudicating, for example, whether a defendant is guilty of a crime or whether a defendant is at fault for a plaintiff's injuries. In the course of resolving a particular dispute, a judicial opinion also has long-term effects. Legal opinions provide rules by which courts should resolve similar disputes in the future (Aldisert 1989), thereby influencing the outcomes of future cases.

While the parties care deeply about the outcome of their dispute, the broader societal effect of a judicial ruling lies in its longer-term effects on the development of law. Existing case law may place constraints on how judges interpret the law, but each case is also an opportunity to craft an opinion that may bind future judges to a particular understanding of the law (Hansford and Spriggs 2006; Bailey and Maltzman 2011). In order to understand how the law develops over time, both in terms of legal constraints and the opportunities available to judges seeking to promote their policy preferences, we must understand what factors make some opinions influential and others obscure.

Despite the evident variation in the long-term effect of judicial opinions and the importance of understanding legal development, scholars have developed and tested relatively few theories to explain this process. Moreover, the existing work tends to focus on opinions generated at the peak of the legal hierarchy (e.g. Hansford and Spriggs 2006; Wahlbeck 1997). In this paper, we develop a more general theory of how opinion characteristics play a role in the complicated process of legal development. We hypothesize that both efficiency and persuasiveness structure future judges' choice of citations and, by extension, the long-term effect that an opinion has on the development of the law. Opinions that are well-grounded in the law and written by authors with expertise have greater persuasive value while precedents that are easier to read and contain fewer footnotes allow future judges to produce their own opinions more efficiently. In short, we expect opinions that can be understood more efficiently and opinions that are more persuasive to have a greater effect on legal development. Undoubtedly, other factors such as case salience, legal relevance, and ideology

also shape legal development. However, this project is focused specifically on the role of opinion characteristics that are subject to judicial manipulation. Our goal is not to describe all determinants of legal impact, but rather to explore the extent to which judges control their opinion's destiny.

We conceptualize legal development in terms of whether and how an opinion is used in future cases. This approach makes it important to account for variation in citation practices across institutional contexts. The principle of stare decisis mandates that a subordinate court cite a legally relevant opinion yet imposes no such obligation on a superior court (Schauer 2008). Furthermore, opinions written in the same court are binding under the doctrine of stare decisis while opinions written by a sister court are not legally binding (Hinkle 2015; Caldeira 1985). To account for these variations, we test our hypotheses separately for each type of citation. In order to be able to test both binding and persuasive citations in both vertical and horizontal contexts, we examine opinions written by courts that are only subordinate to the U.S. Supreme Court.

We test our theory using a core dataset of published search and seizure cases from federal circuit and state high courts from 2000 to 2010 as well as data on every citation to each of those opinions (from any court) through the end of 2013. We estimate the effect of an opinion's readability, the number of footnotes it contains, the number of non-binding precedents it cites, and the general or subject expertise of its author on the effect the opinion has on legal development. We measure this impact in terms of the total number of times an opinion is discussed, the number of times an opinion is treated positively, and the number of times an opinion is treated negatively. There is evidence that efficiency, persuasiveness, and institutional context shape a judicial opinion's legacy.

Overall, our findings make two contributions to understanding how law develops. The primary contribution is that judges have some degree of control over an opinion's future legal impact. Both an opinion's efficiency and persuasiveness can be directly subject to manipulation. A second contribution is also woven throughout our findings: the degree to

which opinion characteristics matter is conditioned by institutional context. For example, the U.S. Supreme Court's extensive agenda control leaves little reason to prioritize efficiency, and we find no significant relationship for such variables in the Supreme Court context. These institutional differences underscore the need for future research that examines legal development throughout the judicial hierarchy. More broadly, our results shed light on the effect that text has on contexts in which law is binding as well as those contexts in which the law is not binding.

Background

Most studies of judicial decisionmaking examine either the outcomes of cases or the votes of judges (Friedman 2006; Knight 2009). Yet the outcome of a judicial decision (and the votes that underlie it) is only a small part of a judicial opinion; judicial decisions must justify those outcomes based on the arguments presented by the parties and the dictates of past cases.

To this end, a bevy of recent research has expanded the judicial behavior project beyond a single-minded focus on case outcomes toward a broader understanding of judicial opinion writing, seeking to understand why judges choose to cite some opinions but not others (Hansford and Spriggs 2006). Legal justifications provide the most obvious explanation for variation in citation practices because judges' justifications are shaped by the doctrine of stare decisis. Judicial opinions on the same topic issued by higher courts and by the same court in years past are binding on a judge's decision in the present case and must be followed. Opinions written by sister courts—those courts who share a rung of the judicial hierarchy with the court deciding the present case—are merely persuasive, meaning that they can provide some support for an opinion but need not be followed (Schauer 2008). Opinions by lower courts are similarly nonbinding. There is empirical evidence that these distinctions set forth by legal doctrine influence how judges use precedent. Hinkle (2015)

finds that federal circuit judges' ideology plays a larger role in the decision to negatively cite persuasive precedents compared to binding precedents.

In addition to decisions about which opinions they should cite, judges face a host of rhetorical decisions when crafting an opinion. Should their opinion be long, featuring an extended discussion of relevant law, or should it be concise? Should the opinion be written technically—accessible only to legal specialists—or written in clear prose? These are all decisions that relate to a judge's available resources: clear, concise writing takes time. Judges, like all professionals, have finite resources. Even at the U.S. Supreme Court, workload and timing considerations play an important role throughout the opinion writing process (Maltzman, Spriggs and Wahlbeck 2000). The justices of the nation's highest court are particular with regard to the sources of the arguments they employ, being more likely to crib language from the lower court opinion when it is a published one written by a prestigious judge and from the parties' briefs when the attorney is experienced and ideologically compatible (Corley 2008; Corley, Collins and Calvin 2011).

Beyond the sources of the legal justifications provided in opinions, other variation in the language employed by judges has a large systematic component. Owens and Wedeking (2011) demonstrate that, while the clarity of a judicial opinion is unrelated to the author's ideology, both dissents and opinions joined by a minimum winning coalition are particularly clear. Similarly, Owens, Wedeking and Wohlfarth (2013) find that the U.S. Supreme Court alters the language of its opinions to evade review by Congress, becoming more likely to obfuscate in its opinions in the face of an ideologically distant Congress. Hinkle et al. (2012) find that U.S. District Court judges engage in similar behavior, becoming more likely to use hedging language when they are not aligned with the majority of the judges on the appellate court that will potentially review their decision.

State supreme court judges are similarly strategic in the language they employ. State supreme court judges are strategic in their opinion writing as it concerns their continued service on the bench, grounding their decisions in state law (thereby making the decisions

non-reviewable by a higher court) when those opinions are likely to cause them electoral problems (Beavers and Walz 1998). Elected judges, at least in recent years, write more readable opinions than their counterparts in states that do not require judicial elections (Nelson 2014).

Less research explores the consequences of these linguistic decisions on the legacy of an opinion. Corley and Wedeking (2014) find that the level of certainty expressed in a judicial opinion affects the treatment it will receive by future courts. Opinions expressing more certainty are more likely to be positively treated by lower courts.

Though this research underscores the strategic nature of opinion drafting throughout the judicial hierarchy, these individual studies tend to only examine the behavior of judges or the citation practices in a single court. Yet, an opinion can be used by any court. The literature is largely silent on cross-court citation practices, particularly as they concern judicial federalism; studies typically examine only a single level of court or citation practices in the federal judicial hierarchy (Caldeira 1985; Hansford and Spriggs 2006). Moreover, extant research focuses primarily on the U.S. Supreme Court or state supreme courts, offering little opportunity to examine the effects of vertical persuasive precedent. We extend this line of work by investigating the consequences of the language used in a judicial opinion both up and down judicial hierarchies.

Efficiency, Persuasiveness, and Hierarchy

Both federal and state judges who answer only to the Supreme Court of the United States have far-reaching effects on the development of law. The doctrine of stare decisis places relatively little constraint on these judges while rendering their decisions decisive throughout their respective jurisdictions. Yet there is still much to learn about the effect of individual precedents from such courts. The ever-increasing proliferation of available precedents complicates the apparently clear doctrine of stare decisis. Courts may not be able to feasibly

address all binding precedents, and the variety of options may present an opportunity to use preferred precedents selectively (Niblett 2010). Moreover, judges from other jurisdictions may choose to ignore, gloss over, or discuss a non-binding precedent for a variety of reasons. Precedents that are discussed more widely have a greater effect on the development of law, but scholars understand fairly little about the factors that determine the effectiveness of a given judicial opinion.

One of the challenges in studying the course of legal development is quantifying the effect a legal decision has on changing the law. Theoretically, a case may have such a tremendous effect by clarifying the contours of the law that it virtually eliminates subsequent litigation, thereby appearing to have little future relevance. While this is perhaps an extreme scenario, it illustrates that the effects of a judicial opinion often range beyond courtrooms. With no practical way to assess such effects, we turn our attention to the development of legal doctrine as it manifests within the legal system. Specifically, we look at the depth and direction of citations to an opinion to gauge its effect on the course of law. Citations to an opinion provide evidence that its reasoning and arguments live on and continue to influence subsequent judges (Cross et al. 2010; Fowler et al. 2007). However, not all citations indicate influence to an equal extent. Mere inclusion in a string citation does not necessarily constitute an impact on legal development in any meaningful sense. On the other hand, actual discussion of a precedent does indicate that the content of the precedent is playing a continued role in the development of law.[1]

The depth of a citation, i.e., whether it is discussed or merely cited in passing, is only one dimension of a citation. Its direction matters as well. A citation can be explicitly directed at expanding or contracting the scope of a precedent (Spriggs and Hansford 2000). A citation constitutes negative treatment when a precedent is addressed for the expressed purpose of limiting, narrowing, or even rejecting its analysis. A positive treatment is a reference to a

[1]We operationalize "discussion" of a precedent as any citation that references multiple points of law from the precedent. See below for further detail.

precedent that explicitly expands the application of legal doctrine set forth therein to the facts of the citing case. While the substantial majority of citations do not reach either extreme, both positive and negative treatments are important manifestations of the extent to which an opinion influences legal development (Hansford and Spriggs 2006). A precedent with more positive treatments has a larger impact on the law because its application has been expanded, either geographically when positively treated in other jurisdictions or to apply to a wider set of facts when positively treated within the precedent's own jurisdiction. Conversely, more frequent negative treatment decreases the influence of a precedent by similarly restricting its scope either geographically or factually.

We now examine how features of an opinion and institutional context can explain variation in legal impact. Potentially relevant cases are brought to a judge's attention by a variety of sources. The litigants on both sides point out precedents in their briefs, lower court opinions contain legal analysis of relevant cases, and a judge's law clerks or other professional staff may provide additional citations. While the judge has access to the relevant choice set through such sources, they are not all similarly available to researchers. Lower court opinions and briefs are not always readily available, and clerks' bench memos and draft opinions are confidential by their very nature. However, the subset of the choice set that is observable is the list of all precedents the judge actually cites in the opinion. The practice of using string citations to list a variety of case law that supports a particular point enables judges to pack a quantity of citations into an opinion. Once describing a particular legal point, it takes very little additional effort to include five cases in the string cite rather than just three. Yet time constraints do play a role in the actual discussion or treatment of precedent. Engaging with a precedent and describing its application to, and implications for, the case at bar take time and effort. Justifying the expansion or contraction of a precedent similarly involves non-trivial effort. As a result, the process of crafting an opinion does not necessarily force choices about which cases from the choice set to cite, but it does typically force choices about which of those cited cases to discuss and treat positively or negatively.

How do judges decide which precedents they cite should receive a discussion or treatment rather than simply a quick citation? More specifically, what characteristics of the precedents themselves play a role in how frequently they are discussed and the extent to which they are expanded or contracted? The substantive importance of these questions are increasing over time as the ever-increasing body of case law leads to larger choice sets (Niblett 2010). The amount of time judges have available to address precedents is fairly constant while the number of relevant cases continues to grow. For example, after the passage of time a set of nine similar case can grow to twelve, and a judge dealing with that issue will still have approximately the same amount of time to discuss (and perhaps expand or narrow) a selection of those cases. While all twelve cases can be included in string citations, it is not necessarily feasible to expand the number of discussed precedents. Instead, judges have to become increasingly selective as the number of available options continues to grow. In this competitive environment, how can a judge craft an opinion to optimize its continued relevance and encourage future use?

We focus on two general factors we expect to influence which types of precedents are discussed and treated more frequently: efficiency and persuasiveness. First, efficiency is likely to play an important role given time constraints. Many judges throughout the court system face a combination of increasing caseloads and limited time and resources with which to resolve those cases (Galanter 2004; Resnik 1982). Precedents written in a manner that makes them more efficient to process with less cognitive effort may be discussed and treated more simply because of their enhanced accessibility. A judge may quickly relegate a convoluted or complex precedent to a string citation regarding a general point while being drawn to engage a carefully crafted precedent in a more nuanced discussion. A judge need not (and often cannot) discuss every applicable precedent, but is simply looking to discuss a subset of relevant opinions sufficient to explicate her rationale and conclusions. Since judges have considerable discretion, busy judges will be likely to exercise that discretion in a way that saves them time. They can do so by choosing to discuss opinions that can be read and

understood in a nuanced way in less time.

While the logic of efficiency is most obvious with respect to which cited precedents are discussed, the same concerns may also impact treatment decisions. A judge may be willing to positively (or negatively) treat a set of precedents that is larger than can be feasibly addressed given time constraints. Under such circumstances, a subset of these opinions must be selected for treatment. Among other factors, judges seeking to reduce the time needed to draft an opinion can pursue that goal by focusing their efforts on expanding or contracting the precedents that can be cognitively processed most quickly. In the case of negative treatment, this means that efficiency concerns may lead to the undesirable effect of a better written opinion opening the possibility of more frequent negative treatment compared to a similar opinion that is drafted in a more convoluted way.

Two measurable features of an opinion make it more efficient for other judges throughout the legal system to swiftly understand its content. The first is simply the opinion's readability. Automated measures of the complexity or accessibility of a text have been available for some time (Flesch 1948). Text with more complex syntax and content is classified as less readable while clear, straightforward, and simple prose is classified as more readable. Surely all judges are capable of reading any opinion, no matter how complex. The key is that more readable texts take less cognitive capacity (and, by extension, less time) to process (Britton et al. 1982). Consequently, we expect a judge scrolling through a series of opinions to choose to discuss and treat opinions that take less time and effort to read. It is necessary to examine at least an excerpt of a case to assess its readability. However, any person who has graded a stack of papers or essay exams knows that it takes surprisingly little time to develop a clear idea of the readability of a particular piece of writing. We hypothesize that more readable opinions are discussed and treated more frequently than less readable opinions.

> *Efficiency Hypothesis 1: More readable opinions are discussed, positively treated, and negatively treated more frequently.*

Readability scores do not account for the number of footnotes. There is considerable

variability among judges in the extent to which they use footnotes.[2] Such variation appears to be somewhat idiosyncratic, and the *Bluebook* has not formally weighed in on the debate (Garner 2014). Regardless of what drives the choice to employ footnotes, that usage may impact how a precedent is used by subsequent judges. Footnotes may play a substantial role in the time necessary to read and process an opinion. As Posner (2001) puts it: footnotes "force the reader to interrupt the reading of the text with glances down to the bottom of the page. They prevent continuous reading. In doing so they make the reader work harder for the same information." (24). This is akin to the difference between driving a mile down a single street and driving the same distance, but doing so by diverging from the main street to drive down a series of dead end streets and then returning to the main road. Even if the same distance is covered, the latter takes more time. While a reader may choose to simply ignore footnotes, even a quick read often involves at least glancing at footnotes to see whether they merit attention. This break in concentration decreases the efficiency of the cognitive processing of the text.

Furthermore, in the context of legal writing, a large number of footnotes may also be seen as an indicator of a writing style that requires substantial effort on the part of the reader. Extensive use of footnotes is a hallmark of law review articles. Judges, all too familiar with that less-than-concise genre, may be less willing to devote the time necessary to discuss a precedent written in that style. For these reasons we hypothesize that precedents with more footnotes will be discussed and treated less frequently.

> *Efficiency Hypothesis 2: Opinions with more footnotes are discussed, positively treated, and negatively treated less frequently.*

[2]The two poles of the debate are two paragons of legal writing. Eminent legal writing expert Bryan Garner is known (e.g. Garner 2014) for his dictate that judges should use footnotes liberally while Judge Posner (e.g. Posner 2001) has long been a vocal opponent of the use of footnotes.

A second general feature of a precedent that influences how much it shapes legal development is its persuasiveness. When faced with a collection of relevant caselaw, a judge will naturally gravitate towards discussing those opinions that contain the more persuasive legal analysis. The dynamic context of legal reasoning means that a judge's own opinion will be viewed as more persuasive if it is grounded in a thorough discussion of particularly persuasive precedents. Put differently, more persuasive precedents provide a stronger foundation which gives subsequent judges more reason to select them for discussion rather than leave them to languish in string citations.

Furthermore, we expect persuasiveness to influence whether a precedent's scope is expanded or narrowed. The persuasiveness of an opinion can influence treatment either directly or indirectly. Direct persuasion occurs when the legal reasoning provided genuinely convinces later courts of the wisdom of expanding the scope of the precedent. Conversely, weakly justified opinions are less convincing, leading directly to more negative treatments narrowing their scope. The persuasiveness of a precedent can also have an indirect effect when judges anticipate how their own opinions will be viewed by others. Similar to the decision about which cited cases to discuss, expanding the scope of more persuasive precedents is more likely to be perceived by outside actors as legally sound. Along the same lines, narrowing a precedent that is less persuasive is less likely to be questioned. While it is difficult, if not impossible, to disentangle direct and indirect persuasion, they both lead to the same theoretical expectations. Therefore, due to the combination of direct and indirect persuasion, we expect more persuasive precedents to be positively treated more frequently and negatively treated less frequently.

Many things that are difficult or impossible to measure contribute to a precedent's persuasiveness. To address this challenge, we examine measurable features of a precedent and its author that may contribute to an opinion's actual or perceived persuasive value. First, we consider the use of precedent within an opinion. Given the central role of stare decisis in a common law system, opinions that are well-grounded in existing precedents are more

likely to be perceived as well-reasoned and persuasive by subsequent judges (Hansford and Spriggs 2006). The judicial task necessarily requires addressing relevant binding precedents, but a judge is not required by legal doctrine to extend the analysis to also consider relevant caselaw from other jurisdictions or from lower courts. Dealing with such non-binding precedents is entirely up to the discretion of an authoring judge. An opinion supported by reference to, and discussion of, a wide swath of non-binding precedents is more persuasive because its legal reasoning is supported by a wider range of outside sources. Such citations also provide evidence of a judge investing more effort in strengthening an opinion's reasoning. As a result, opinions that contain more citations to non-binding precedents are potentially more persuasive than opinions with fewer such citations. We hypothesize that precedents with a greater number of citations to non-binding caselaw will have a greater influence on the development of law because they will be discussed and positively treated more frequently and negatively treated less frequently.

> *Persuasiveness Hypothesis 1: Opinions that contain more citations to non-binding case law are discussed and positively treated more frequently and negatively treated less frequently.*

The second indicator of persuasiveness we consider is how the expertise of an opinion's author may serve as a proxy for the ultimate persuasive value of the opinion itself. Just as persuasion can operate directly or indirectly, so expertise may contribute to persuasive value directly or indirectly. First, expertise might make a judge more capable of writing higher quality, more persuasive opinions. For example, previous professional experience relevant to a case or an education at a more prestigious law school may better equip a judge to craft particularly persuasive reasoning. Second, expertise might strengthen a judge's reputation leading to the perception that her opinions provide a stronger, more convincing, analytical foundation that opinions written by other judges. These elements are heavily intertwined, but they have the same implications; author expertise should enhance an opinion's legal impact. While it is difficult to measure expertise, attending an elite law school

is a widely-acknowledged, if blunt, proxy for the ability to craft exceptionally well-reasoned legal reasoning. Previous career experience relevant to a particular legal topic is a signal of subject-area expertise. We hypothesize that both general and subject-specific expertise of the author should lead to an opinion being discussed and expanded more and narrowed less.

> *Persuasiveness Hypothesis 2: Opinions written by judges with greater expertise are discussed and positively treated more frequently and negatively treated less frequently.*

The first four hypotheses provide our general expectations for how persuasiveness and efficiency influence a precedent's impact on legal development. The hierarchy of the judicial system creates institutional differences that lead us to expect further variation in how precedents are used. The doctrine of stare decisis makes decisions of higher courts binding on lower courts and within the same court at a later date, but across jurisdictional lines precedents are only persuasive (Hinkle 2015). To parse the effects of institutional structure, we unpack our examination of all citations by looking at four types of potential citing courts categorized based on their relationship to the court that crafted an opinion. Figure 1 illustrates this breakdown of the potential sources of citation. First, a precedent may be cited by the U.S. Supreme Court, the direct superior of the opinion-writing court. Second, an opinion may be cited by the same court at a later date. Third, an opinion may be cited by a sister court, defined as a different court at the same level. For federal circuit cases, the other federal circuits constitute sister courts. For state courts of last resort, the analogous courts from other states constitute sister courts. Fourth, another possibility is that an opinion may be cited by the lower courts that are directly subordinate to the opinion-writing court.[3]

[3]Although the phrase "lower courts" may be used more generically to simply refer to trial courts, in this paper we use the phrase to specifically denote those trial courts that are in the same jurisdiction as the opinion-writing court.

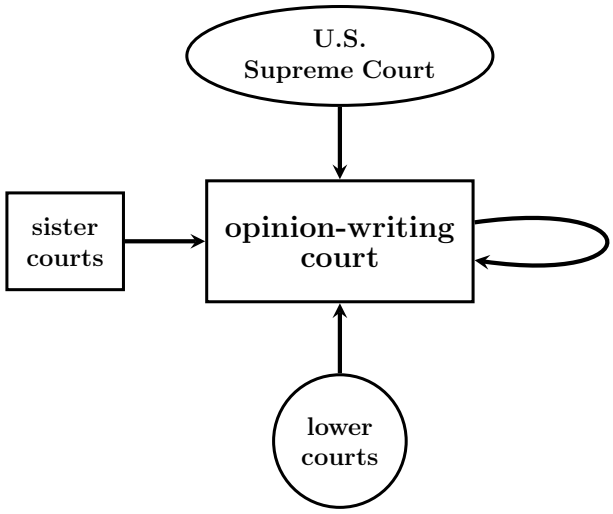

Figure 1: Potential Sources of Citation to A Federal Circuit or State High Court Opinion

The relevance of efficiency and persuasiveness is not necessarily universal across institutional contexts. The importance of time and resource constraints means that efficiency is most likely to have an effect where those concerns are paramount. In the context of the U.S. Supreme Court, this logic is least applicable. The justices have both sufficient time to look at each case in considerable depth and greater resources to assist that process than any other court in the land. Moreover, the Court's important role making legal policy for the entire country means that accuracy and thoroughness are prioritized much more highly than efficiency. At the other end of the spectrum, lower courts face the largest caseloads and have considerably less time to devote to resolving each individual case than their superior courts do. As a result, this context is one where efficiency is most particularly important. Due to these differences in the importance of saving time and resolving cases quickly, we expect efficiency to have the smallest effect on how the U.S. Supreme Court uses precedents and the largest effect on how lower courts use precedents.

> *Efficiency Hypothesis 3: Opinion readability and the number of footnotes will have the smallest effect on the use of precedent by the U.S. Supreme Court and the largest effect on the use of precedent by lower courts.*

Like efficiency, the importance of an opinion's persuasiveness should depend on context. First, persuasiveness is not likely to influence the U.S. Supreme Court very much. Given their resources and unique position at the apex of the legal hierarchy, Supreme Court justices have both the luxury and responsibility to assess every legal issue in nuanced and complex detail. Furthermore, justices do not need to seek out particularly solid lower court precedents in order ensure capturing the attention of future judges. Second, persuasiveness should also matter less in contexts where a precedent is legally binding under the doctrine of stare decisis. Lower courts and the same court the opinion came from are potentially constrained by legal doctrine (Aldisert 1989). There may still be room for persuasion to impact how a precedent is used if a multitude of available binding precedents offer a choice. As a result, persuasion should play the largest role in how an opinion is used by sister courts. But the role an opinion's persuasiveness plays in how much it impacts legal development should be dampened when the precedent is binding compared to when it is not binding. This hypothesis and all the forgoing hypotheses are summarized in Table 1.

> *Persuasiveness Hypothesis 3: The number of non-binding precedents cited and author expertise will have a smaller effect on the use of precedent by the U.S. Supreme Court, lower courts, and the same court and the largest effect on the use of precedent by sister courts.*

	All Cts.	SCOTUS	Same Ct.	Sister Cts.	Lower Cts.
Discussion & Positive Trt.					
Efficiency					
Readability	+	+	+	+	+
Footnotes	-	-	-	-	-
Persuasiveness					
Non-Binding Cases Cited	+	+	+	+	+
Elite Law School	+	+	+	+	+
Subject Expertise	+	+	+	+	+
Negative Treatment					
Efficiency					
Readability	+	+	+	+	+
Footnotes	-	-	-	-	-
Persuasiveness					
Non-Binding Cases Cited	-	-	-	-	-
Elite Law School	-	-	-	-	-
Subject Expertise	-	-	-	-	-

Table 1: Summary of Hypotheses

Data and Research Design

Not all legal issues provide the same level of opportunity to influence legal development. Some more obscure areas of law may lead to less frequent citation of an opinion simply because the topic is rarely litigated. Similarly, if one looked for empirical patterns across the full range of judicial decisions, one would need to account for differences in the legal content in each case. For these reasons, we focus our empirical analysis on one issue area: Fourth Amendment search and seizure law. This topic is well-suited for this study because it incorporates a discrete set of legal issues that are routinely raised in both state and federal litigation, and relevant cases can be identified by the simple expedient of finding cases that cite the constitutional search and seizure clauses which bind the court (the U.S. Constitution for all courts and the relevant state constitutional provision for each state supreme court).[4] Using Lexis, we collected every such published opinion from a federal circuit court or a state

[4]A legal publication, *Shepard's Citations*, provides this list.

court of last resort between 2000 and 2010.[5] After excluding all opinions that do not address the merits, do not identify the author, or are too brief for textual analysis (i.e., less than 150 words), the resulting dataset contains 7,503 cases.

For each opinion in our dataset, we collect information from *Shepard's Citations* on every time a court (of any type) cited that precedent up through December 31, 2013.[6] Our core dataset of search and seizure opinions generated a total of 676,011 citations during this time frame. In order to evaluate the effects of institutional structure, we identify which citations are in each of the four categories discussed above: citations from the U.S. Supreme Court, the opinion-writing court, sister courts, and lower courts. Figure 2 shows the breakdown of citations into these four categories. These four types of citations constitute 70% of all citations in our data. The remaining 30% of citations cross both jurisdictional and hierarchical boundaries in a variety of ways. Although these miscellaneous citations are included in the overall totals for each opinion, we do not examine these as a separate category due to the considerable heterogeneity in context. This collection of over half a million citations constitutes our primary data source. Using the information provided in *Shepard's* we classify each citation based on whether it discusses, expands, or narrows the opinion in question.

Our unit of analysis is the opinion. We examine three sets of outcome variables for each opinion. The first is the number of times an opinion has been discussed as of December 31, 2013. Even the most recent opinion in our dataset has been available for discussion for three years. Most citation research utilizing the extensive information compiled by *Shepard's Reports* simply relies on the distinction between citation and the treatment categories used by *Shepard's*. While these sources are valuable (and we use them as well) there is a substantial

[5]Federal appellate courts with subject-specific jurisdiction are not included in this study.

[6]All discussion of citations is limited to citations to the majority opinion in a case. *Shepard's* also contains information on citations to separate opinions. These are excluded from all analysis.

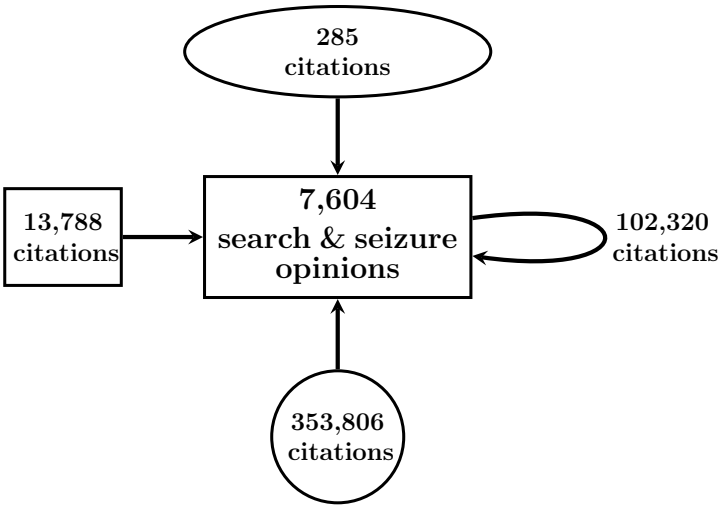

Figure 2: Summary of Citations by Type

amount of heterogeneity remaining among citations that do not earn a directional treatment designation (either positive or negative) under the *Shepard's* coding scheme. A citation that does not rise to the level of a positive or negative treatment may still contain a meaningful discussion of the precedent in question. Or it may be no more than one reference in a string cite used to support the most straightforward and uninteresting of legal principles (e.g., the applicable standard of review).

We move the literature on legal precedent forward by developing a measure that distinguishes between mere citation and actual discussion of a precedent. One of the pieces of information included in *Shepard's* about a citation is which Lexis Headnote(s) from the precedent are referenced in the citing opinion.[7] If a citation addresses more than one headnote from a precedent, that indicates the citing case is talking about multiple legal points from the precedent. We leverage this feature by collecting a count for each citation of the number of headnotes addressed. Any citation that addresses two or more headnotes (and is not a positive or negative treatment) is coded as a "Discussion" of the precedent. While necessarily blunt, this approach provides a way to classify over half a million citations that still provides a meaningful separation between citations that may simply be string cites and

[7]Headnotes are the summary of legal principles Lexis assigns to each opinion.

those that plausibly contain a more nuanced look at the precedent in question. Just over 20% of citations constitute a discussion of precedent using our approach. This is similar to the 19% of citations with *Shepard's* treatment categories indicating positive or negative treatment.

The second and third type of outcome variable we track capture the directionality of a citation. Positive treatments are citations that expand the scope of a precedent, while negative treatments narrow or restrict its potential application. Some treatment categories in *Shepard's* can be both ambiguous and heterogeneous, so we follow the advice of Spriggs and Hansford (2000) and only utilize treatment categories that clearly indicate either positive or negative treatment. Using their classification, 'Followed' is the only *Shepard's* treatment classified as positive while negative treatments include 'Distinguished,' 'Criticized,' 'Limited,' 'Questioned,' 'Overruled,' and 'Disapproved.' All three types of outcome variables are counts that show evidence of overdispersion, so we model them using negative binomial models. For each set of outcome variables, we estimate separate models to examine the use of a precedent overall and by the four different types of courts. Since precedents issued by different courts and at different times may have varying impact, we estimate robust standard errors clustered on the specific court and year in which the precedent was written.

The key explanatory variables are *Readability*, *Footnotes*, *Non-Binding Cases Cited*, *Elite Law School* and *Subject Expertise*. Our measure of *Readability* uses factor analysis to combine six different common measures of opinion readability into a single measure; Hansford and Coe (2014) provide some experimental evidence that these measures are valid when applied to judicial opinions. More extensively, Black et al. (2016) validate automated measures of readability using human raters, presenting substantial evidence that raters have similar perceptions of opinions classified by automated measures as less readable.[8] More information

[8]Specifically, Black et al. (2016) demonstrate that human judgments of an opinion's readability, writing quality, comprehensibility, vocabulary, ease of reading comprehension, and the amount of time taken to read an excerpt are all highly reliable (alpha=0.77) measures

about the readability scores used to formulate the measure is available in Appendix A. The variable has mean zero and a standard deviation of one with higher scores indicating more readable judicial opinions. *Footnotes* is the number of footnotes in a majority opinion. This measure was extracted using a Python script with follow-up checks to verify accuracy. In addition to information on cases which cite an opinion, *Shepard's Reports* also provides a list of all case law cited within an opinion. Using the Table of Authorities for each opinion in our dataset, we extracted the total number of non-binding judicial opinions cited by an opinion for the variable *Non-Binding Cases Cited*. Finally, *Elite Law School* is a dichotomous variable that equals one when the author of an opinion went to a top 14 law school[9] and zero otherwise, and *Subject Expertise* equals one when the authoring judge has prior experience as a prosecutor and zero otherwise.

As discussed in the previous section, we control for characteristics of the opinion, its author, and the originating court that might be correlated with both the use of precedent and our key explanatory variables. First, we control for the total number of citations to a precedent in order to account for differences driven by variation in case salience, number of legal issues addressed by an opinion, and frequency of litigation on the particular legal questions an opinion addresses. For similar reasons, we further control for the length of an opinion. At first blush, the length of an opinion might appear to be a useful proxy for efficiency. However, a major component of opinion length is not merely the author's prolixity, but simply the number of legal issues that need to be addressed. Since longer

of the concept. Regressing the rater's comprehension of an opinion on the automated measure of readability, Black et al. (2016) conclude "excerpts that our automated readability statistics identified as being easier and more readable yielded higher comprehension levels than excerpts that were more difficult. These findings provide systematic support for our computer generated readability measure" (76).

[9]The top-14 law schools are (in alphabetical order) Berkeley, Chicago, Columbia, Cornell, Duke, Georgetown, Harvard, Michigan, NYU, Northwestern, Penn, Stanford, UVA, and Yale.

opinions tend to address a wider variety of legal topics, they can also be expected to be discussed or treated more frequently than shorter opinions. Therefore, we control for the natural log of the number of words in an opinion.

Opinions that must address a greater number of issues may also tend to cite more non-binding precedents without generating an opinion that is, in fact, more persuasive. In order to isolate the persuasive effect of using non-binding precedents, we control for the number of binding precedents in an opinion. A more complex case that requires the author to resolve more legal nuances or a case that involves issues with fewer available binding precedents may affect how many non-binding precedents are cited and such variations are not necessarily linked to the persuasive value of an opinion. However, when comparing two opinions with the same number of citations to binding precedent, the opinion with more extensive use of non-binding precedents should provide greater persuasive force. The count of binding precedents cited in each opinion is obtained from *Shepard's* Table of Authorities for each case just like the count of non-binding precedents.

Some courts have an institutional reputation that enhances the frequency with which other courts turn to their opinions. In order to account for this, we control for the *Baseline Citation Rate* to a particular court. For each opinion this variable is the total number of citations to the opinion-writing court divided by the total number of citations in the dataset. Importantly, this accounts for any differences in the use of opinions from federal circuit courts and state high courts. The extent of an opinion's effect on the development of law will also be related to the length of time it has been available to be used by other courts. Although our data include opinions written over the course of only one decade, it is still important to account for the difference between an opinion available for 13 years and one available for three years. The variable *Years Available* captures this quantity. Additionally, each type of treatment of an opinion may be related to how often it has received the opposite type of treatment. Consequently, in all treatment models we also control for the total number of times an opinion has been treated the other direction in the applicable type of court.

The one remaining control variable is the ideological extremity of the opinion's author. Since our research design aggregates system-wide use of a particular opinion, whether its author is liberal or conservative is not likely to provide theoretical traction. However, the ideological extremity of an author may very well play a role. Subsequent judges may be less willing to focus on opinions written by judges who are less moderate and more ideologically extreme. In order to construct a measure of such extremity, we begin with two widely-employed measures of judicial ideology; Judicial Common Space scores for federal judges (Epstein et al. 2007) and Bonica and Woodruff's (2012) measure for state high court judges. Taking the absolute value of these respective measures provides a metric of ideological extremity, but they are not on comparable scales for state and federal judges. The left panel in Figure 3 demonstrates that the distribution of raw ideological extremity scores is very different for state and federal judges. In order to place the ideology extremism of state and federal judges on the same scale, we convert both to relative rankings and use the percentile of a judge's ideological extremity compared to all other judges of the same type in the data. For example, Justice Roy Moore of the Supreme Court of Alabama is more ideologically extreme than 92% of all state judges who authored at least one opinion in our dataset. The value of 92 for the variable *Ideological Extremism Percentile* provides a meaningful point of comparison to federal judges, only 8% of whom will have an equal or greater value for this variable. The right panel of Figure 3 illustrates that using the relative rankings of state and federal ideology measures results in much more similar distributions for state and federal judges.

We developed this research design to focus on examining the cumulative impact of a judicial decision. While our approach provides the leverage we seek on the question of what role opinion crafting plays in legal development, one of the limitations is that it requires we disregard the nuances inherent in the differential characteristics of the judges who decide whether and how to use precedent. Although we account for institutional context by looking at citations from different types of courts, many specific characteristics of citing courts and

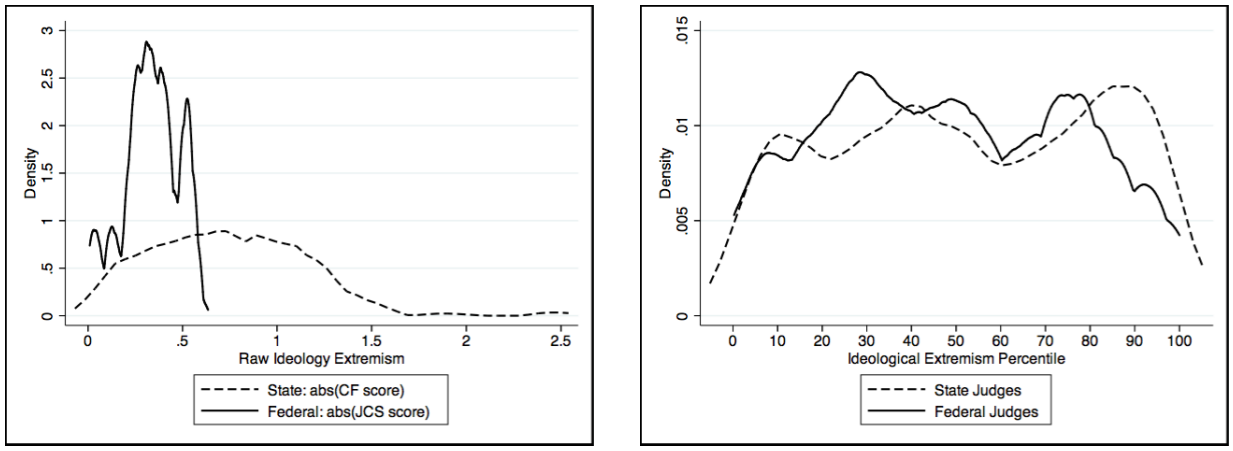

Figure 3: Distribution of Ideological Extremity for State and Federal Authors.

judges are also likely to play a role in citation and treatment decisions. These complexities certainly promise to provide interesting fodder for future work. However, for the purposes of this project we simply note that while a variety of factors drive individual citation decisions, we can usefully examine the aggregate effect of all such decisions across types of courts to learn something useful about the role of precedent characteristics.

Results

Our first four hypotheses concern the overall impact of efficiency and persuasiveness on the use of precedent. In order to test these hypotheses with the broadest set of information available, we begin by modeling the discussion and treatment of our precedents in all courts. Table 2 provides the results of these models. Our first two efficiency hypotheses are supported with only one exception. As we hypothesize, greater *Readability* leads to significantly more discussion and positive treatment of a precedent, but it does not have a statistically significant impact on negative treatment.[10] Consistent with our second efficiency hypothesis, an increase in the number of footnotes in a precedent significantly decreases the frequency with which it is discussed, positively treated, and negatively treated.

[10]All discussion of statistical significance is at the 0.05 level.

The support for our general persuasiveness hypotheses is more limited. An increase in the number of non-binding precedents cited does significantly increase the frequency of discussion as anticipated, but neither of the anticipated effects on treatment emerge. The only support for our expertise hypothesis is the finding that opinions written by judges with prosecutorial experience are less likely to be treated negatively. Alongside these expected findings, there are some surprising results for the persuasiveness variables. Citing more non-binding precedents counter-intuitively increases the frequency of negative treatment, and opinions written by prior prosecutors, who presumably have expertise on search and seizure issues, are discussed and positively treated less frequently. In order to illustrate the substantive size of all these effects, we turn now to an examination of how each outcome varies over the range of each of our measures of efficiency and persuasiveness.

Figure 4 illustrates how efficiency concerns impact the frequency with which a precedent is discussed, positively treated, and negatively treated. As shown in the panel on the left, the least readable opinions in our data are discussed by an estimated 4.9 cases while this number increases to 7 for the most readable opinions in the data. As shown on the right, similarly moving the number of footnotes in an opinion from the minimum to the maximum decreases the predicted number of discussions from 6.2 all the way down to 2.3. Efficiency has a similar size impact on the number of positive treatments with the effect of *Readability* being even more pronounced. Over the range of *Readability* the number of predicted discussions increases from 3.7 all the way to 8.7. The dotted lines in Figure 4 reveal that efficiency plays a much smaller role in the decision to negatively treat a precedent. Negative treatments remain at consistent levels across the range of *Readability* and only decreases slightly as the number of *Footnotes* increases.

Next we turn to examining the substantive effect size of our measures of persuasiveness. Figure 5 graphs the predicted outcomes for *Non-Binding Cases Cited* on the left and for the two different types of author expertise on the right. Moving the number of non-binding cases cited in a precedent from its minimum to maximum increases the predicted number

	Discuss	**Pos. Trt.**	**Neg. Trt.**
Readability	0.040*	0.094*	−0.008
	(0.014)	(0.015)	(0.020)
Footnotes	−0.005*	−0.004*	−0.004*
	(0.001)	(0.001)	(0.002)
Non-Binding Cases Cited	0.004*	−0.002	0.015*
	(0.001)	(0.001)	(0.002)
Elite Law School	0.023	0.025	0.045
	(0.028)	(0.026)	(0.042)
Subject Expertise	−0.100*	−0.080*	−0.136*
	(0.026)	(0.025)	(0.041)
Total Citations Received	0.007*	0.006*	0.001
	(0.000)	(0.000)	(0.000)
Logged Word Count	0.621*	0.494*	0.304*
	(0.036)	(0.030)	(0.060)
Binding Cases Cited	−0.003*	−0.003*	−0.002
	(0.001)	(0.001)	(0.001)
Court Baseline Citation Rate	0.004*	0.003*	0.004*
	(0.000)	(0.001)	(0.001)
Author Ideological Extremity	−0.002*	−0.001	−0.002*
	(0.000)	(0.000)	(0.001)
Years Available	0.048*	0.003	0.073*
	(0.006)	(0.005)	(0.008)
Other Type of Treatment		0.050*	0.016*
		(0.006)	(0.002)
Constant	−4.200*	−2.715*	−3.272*
	(0.303)	(0.245)	(0.503)
N	7503	7503	7503

Table 2: All Courts Models. Negative binomial regression estimates of the effect of opinion and author characteristics on the total number of times an opinion is discussed, positively treated, or negatively treated by a subsequent court. Robust standard errors, clustered on the precedent court and year, are reported in parentheses below each coefficient, and * denotes a p-value less than 0.05.

of discussions from 6 to 10. This change is substantial and indicates that more persuasively crafted opinions are discussed more frequently. Oddly enough, the same opinion characteristic also increases the frequency of negative treatment and it does so to an even greater extent, moving from 1.1 to 9.5 over the range of the variable. However, it is worth noting that the steeper effect of non-binding precedents observed at the higher end of the scale is imprecisely estimated due to the relative scarcity of observations with such values. The right

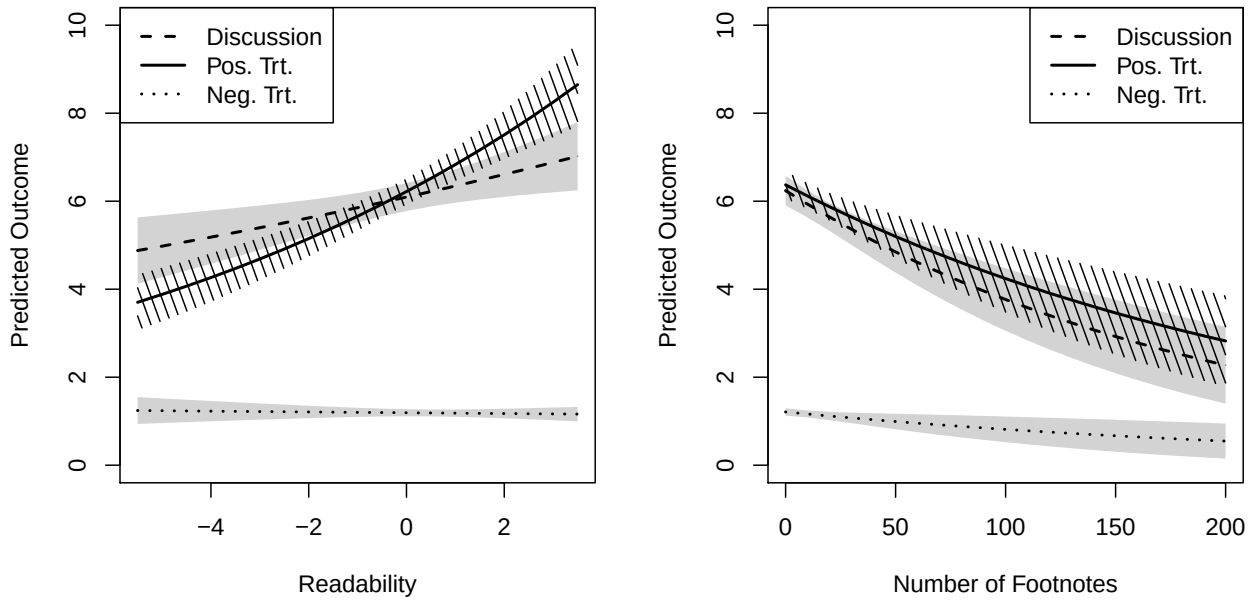

Figure 4: Predicted Outcomes over the entire range of *Readability* and *Footnotes* in turn while holding all other variables at their median. The shaded region around each line denotes the 95% confidence interval.

panel showing the effect size of author expertise highlights the fact that even the statistically significant findings for expertise are quite small in size. This is equally true for the findings in the expected direction and those that are surprising. As we expected, opinions by prosecutors are less likely to be negatively treated, but the empty circle showing the predicted number of negative treatments is only marginally lower than either an author from an elite law school who has never served as a prosecutor or a judge with neither general nor subject expertise. A prosecutor's opinions generate a predicted 1.04 negative treatments while an opinion by a judge with neither type of expertise generates a predicted 1.19 negative treatments. The unexpected findings that prosecutors' opinions are discussed and positively treated less frequently are similarly small in scope.

We are confident the results in Table 2 reflect the impact of efficiency and persuasiveness because the control variables account for potential confounding factors such as issue salience, litigation frequency, institutional reputation, and ideological extremity. As we expect, longer opinions are discussed and treated more frequently, likely because they resolve more legal

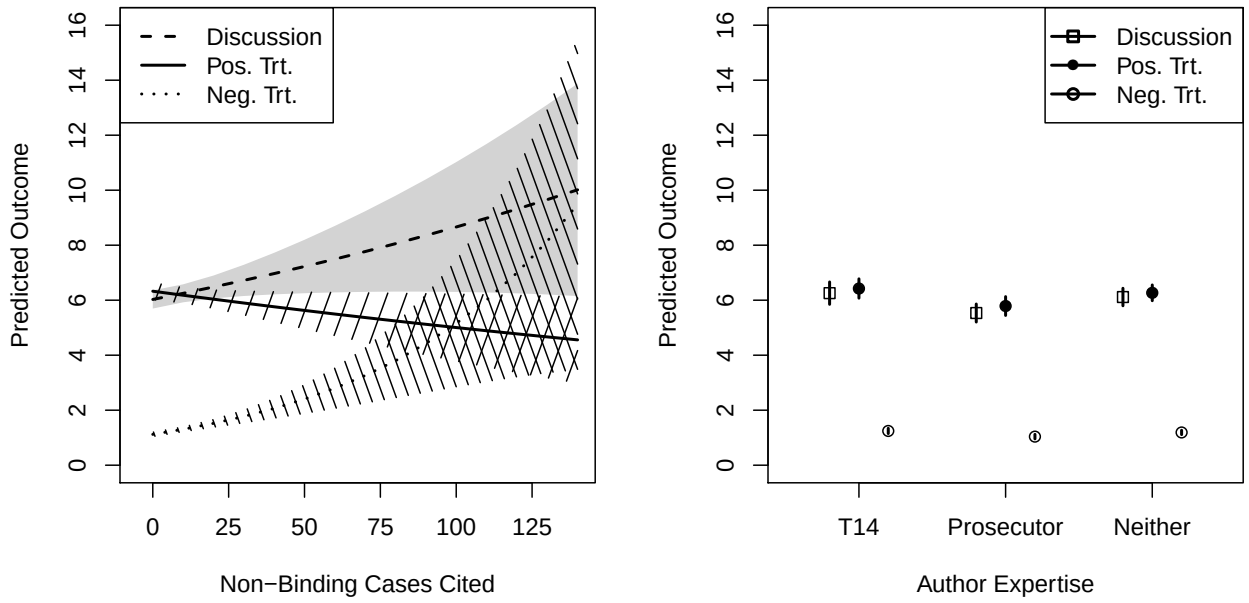

Figure 5: Predicted Outcomes over the entire range of *Non-Binding Cases Cited* and for each type of author expertise while holding all other variables at their median. The shaded region around each line (on the left) and lines through each point estimate (on the right) denote the 95% confidence intervals.

issues than shorter opinions. Opinions from courts that enjoy a higher overall citation rate are also more likely to be treated and discussed. Further controlling for the total number of citations to each particular precedent accounts for differences in treatment and discussion attributable to differences in how frequently a topic is litigated. Finally, we account for the possibility that the known ideology of the author of a precedent may have an impact beyond the content of the opinion. Not surprisingly, precedents written by more ideologically extreme judges are discussed less frequently than those from their more moderate peers. Interestingly, such ideological extremism also reduces the number of negative treatments.

Having examined in some detail the overall effect of efficiency and persuasiveness on how opinions are used by subsequent courts, we now turn to evaluate our hypotheses regarding the conditioning effect of institutional context. We do so by modeling the effect of opinion efficiency and persuasiveness on the number of times an opinion is discussed in each of four specific types of courts: the United States Supreme Court, the same court as the opinion,

sister courts, and lower courts directly subordinate to the opinion-writing court. Simply comparing coefficients across different models does not provide a formal statistical test of whether a variable has a larger or smaller effect in one model compared to another. Therefore, we estimate the change in predicted outcome generated by moving continuous variables from their 25% to 75% values and moving dichotomous variables from 0 to 1. An effect from one model that does not fall within the confidence interval for the same variable in a different model has a significantly different marginal effect. Figure 6 plots the predicted change in discussion for each of the four court types. Additionally, all point estimates with a confidence interval that includes zero are presented in gray (instead of black) to clarify which effects are not statistically significant within their own model.

Efficiency Hypothesis 3 anticipates that *Readability* and *Footnotes* will have the largest effect in the lower courts and the smallest effect in the Supreme Court. Figure 6 illustrates that there is support for this hypothesis. The estimates for *Readability* show that the change in predicted discussion is smallest in the Supreme Court model (as shown by the empty triangle) and largest in the lower court models (as shown by the solid diamond). The effect of *Readability* is significant smaller in the Supreme Court model compared to the Same Court model and the Lower Courts model. Furthermore, the effect of *Readability* is significantly larger in the Lower Courts model compared to the Sister Courts model.

Turning to *Footnotes*, the second measure of efficiency, the results are similar. The magnitude of the (negative) effect of *Footnotes* is also largest in the Lower Courts model and smallest in the Supreme Court model. The magnitude of effect in the Lower Courts model is significantly larger than in the Sister Courts model, but not the Same Courts model. The magnitude of the effect of *Footnotes* in the Supreme Court model is smaller than in all three other models, and it is also not significantly different from zero. While the substantive size of these differences is quite small, they are consistent with our theoretical expectations about the way institutional context conditions the effect of efficiency on discussion of precedent.

Recall that Persuasiveness Hypothesis 3 predicts that measures of persuasiveness will

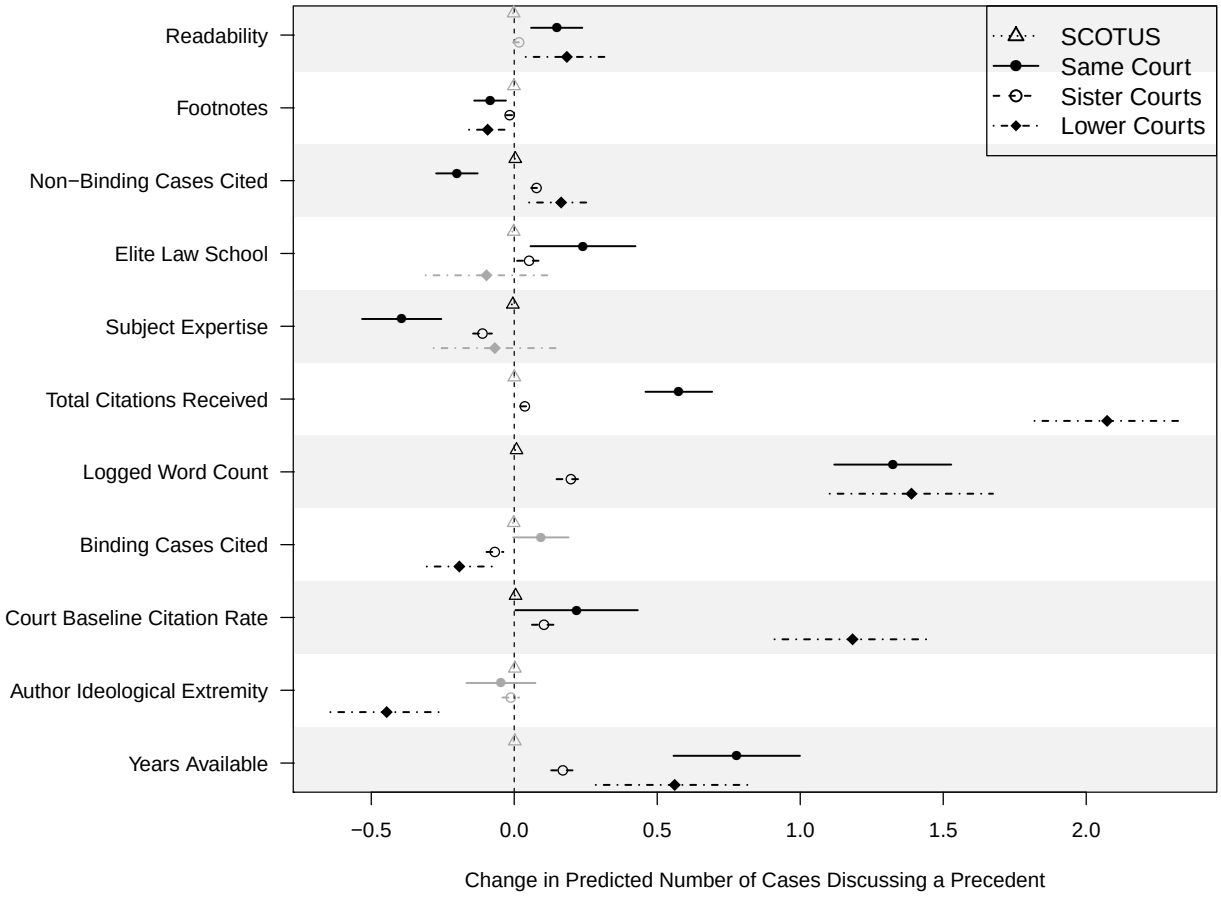

Figure 6: Effect of each variable on the predicted number of cases discussing a precedent when moving continuous variables from their 25% to 75% value and when moving dichotomous variables from zero to one. Full regression estimates are available in the Appendix.

have a larger effect in the Sister Courts model than in any of the other three contexts. In the discussion models there is only sporadic support for this hypothesis. *Non-Binding Cases Cited* has a significantly larger effect in the Sister Courts model compared to the Supreme Court model, but not compared to the other two models. The results for *Elite Law School* are similar. The the third measure of persuasiveness, *Subject Expertise*, only has a significantly larger effect in the Sister Courts model compared to the Same Court model, but it is worth noting that in both instances the effect is unexpectedly negative. Even where the differences are statistically significant and in the hypothesized direction, the substantive size of these effects is quite small.

Finally, we evaluate the conditioning effects of institutional context on how efficiency and persuasiveness influence positive and negative treatment of an opinion. Figure 7 illustrates the changes in predicted outcomes for both positive treatment (on the left) and negative treatment (on the right). The data reveal that both types of treatment by the U.S. Supreme Court are exceedingly rare. In fact, there is too little variation on the outcome variable to identify either model. As a result, we are left with three models for each type of treatment. we expect that efficiency should have the largest impact in lower courts and persuasiveness should most impact on treatment by sister courts.

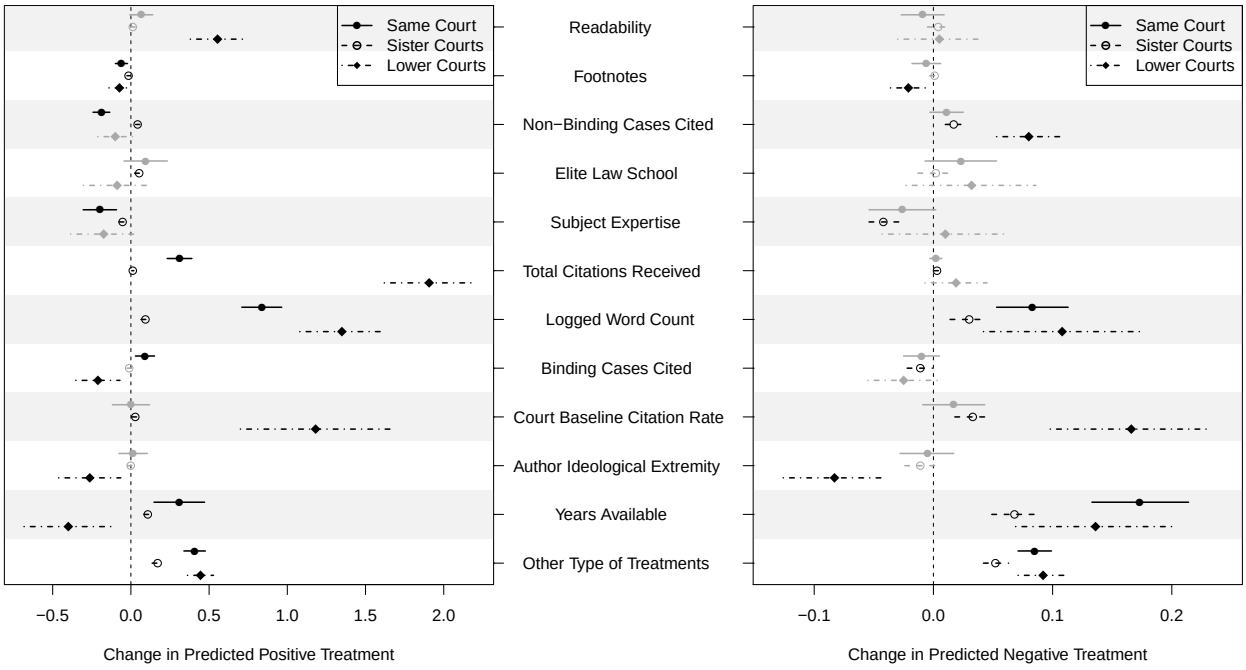

Figure 7: Effect of each variable on the predicted number of positive and negative treatments when moving continuous variables from their 25% to 75% value and when moving dichotomous variables from zero to one. Full regression estimates are available in Appendix B.

There is evidence that efficiency is a greater concern in lower courts, as we hypothesize. *Readability* has a significantly larger impact on positive treatment in lower courts than either sister courts or within the same court. However, the impact of a precedent's readability on negative treatment is statistically indistinguishable among the three models. The results for

the impact of *Footnotes* are more widely consistent with our hypotheses. A greater number of footnotes leads to a smaller number of both positive and negative treatments in the lower court models compared to the other models. Moreover, all of the relevant differences are statistically significant with the exception that *Footnotes* has a similar impact on positive treatment in the lower courts and same court models.

There is only limited support for our expectations regarding how institutional context conditions the effect of persuasiveness. The magnitude of the effect of *Non-Binding Cases Cited* on positive treatment is significantly larger in the Sister Courts model compared to the other two, but there is no similar support for our hypothesis in the negative treatment model. The effect of the author of an opinion attending an elite law school is not significantly different across models for either type of treatment. Finally, the impact of prosecutorial experience on negative treatments is not significantly different across context. The effect of prosecutorial experience on positive treatments is significantly less negative in Sister Courts model compared to the Same Court model, but both are negative.

Discussion and Conclusions

We sought to determine the extent to which an opinion author can determine the subsequent influence her opinions will have on legal development. We argued that two factors—persuasiveness and efficiency—affect the extent to which an an opinion is discussed by subsequent decisions and that the persuasiveness of that opinion affects the extent to which future authors are willing to treat the decision positively or negatively. Furthermore, we sought to track vertical influence both up and down the judicial hierarchy as well as evaluate horizontal influence both within the precedent's own jurisdiction and across jurisdictional lines, suggesting that the extent to which efficiency and persuasiveness affect legal development is conditional upon institutional context. To this end, we provide one of the first analyses to examine legal development throughout the judicial hierarchy.

Our results provide evidence that the legacy of an opinion is influenced both by its own persuasiveness and by later judges' concerns for efficiency. First, we theorized that, because judges have only a finite amount of time, efficiency concerns would predict an opinion's future influence. The empirical evidence generally supports this theory. Opinions that are more readable tend to be discussed and positively treated more frequently while opinions that are laden with footnotes tend to be discussed, positively treated, and negatively treated less frequently as time passes. Although the increase in negative treatments when an opinion has fewer footnotes appears to create a somewhat perverse incentive to craft more complicated opinions, the size of this effect is dwarfed by the countervailing patterns for discussion and positive treatment. On balance, a judge has more to gain in terms of impacting legal development by limiting the number of footnotes.

Second, we hypothesize that, because judges have a desire for good legal policy, opinions that are particularly persuasive, as indicated by how well-grounded the opinion was in extant precedent and the experience of its author, would be particularly influential. Again, there is some evidence for this assertion, especially when one uses the number of non-binding cases cited in an opinion as an indicator of persuasiveness. Authorial expertise, on the other hand, does not have the expected effects on citation, with those judges who have subject-specific expertise (prior service as a prosecutor in our application) actually writing opinions that are *less* influential.

Third, our theory anticipates that institutional context conditions the role of both efficiency and persuasiveness. By casting our net widely to examine use of an opinion in a variety of institutional contexts, we unmasked important variation within the judicial hierarchy. The results demonstrate that citation practices on the U.S. Supreme Court differ markedly from the rest of the court system. We expected that the effect of persuasiveness and efficiency varies throughout the judicial hierarchy. We found support for this hypothesis as it concerns efficiency, with the effect of both measures of efficiency strongest among those courts who deal with the largest number of cases and weakest at the U.S. Supreme Court

whose limited caseload limits its need for efficiency. Again, the results for persuasiveness are more complicated. The effect of persuasiveness is generally stronger among Sister Courts than in the U.S. Supreme Court, as expected. However, the magnitude of the effect of persuasiveness among Sister Courts is indistinguishable from its affects in Lower Courts or in the precedent's own court.

These contextual differences in the effects of persuasiveness and efficiency on legal development underscore the importance of examining judicial politics throughout the judicial hierarchy. While individual levels of the hierarchy can tell part of the story, these results indicate that a single-minded focus on any individual court—particularly the U.S. Supreme Court—may tell a story that is not generalizable throughout the judicial hierarchy. Thus, in addition to the obvious insights provided by the presence of institutional variation in an analysis such as this one, an examination of a fuller judicial hierarchy provides substantial opportunities for additional theoretical development.

There are two caveats to this study. First, focusing on a single issue area does pose some concerns about generalizability. However, there is no reason to suspect that the use of precedent in search and seizure law is substantially different from other areas of law. If anything, that possibility suggests the potential utility of further work investigating variation in legal impact across topic areas. A second caveat is that we do not intend this study to be an exhaustive description of features that influence how much an opinion affects the development of law. Persuasiveness and efficiency are but two factors in what is certainly a highly complex process. We look forward to the challenge of exploring other factors in greater detail in future work.

References

Aldisert, Ruggero J. 1989. "Precedent: What It Is and What It Isn't: When Do We Kiss It and When Do We Kill It." *Pepperdine Law Review* 17(3):605–636.

Bailey, Michael A. and Forrest Maltzman. 2011. *The Constrained Court: Law, Politics, and the Decisions Justices Make*. Princeton University Press.

Beavers, Staci L. and Jeffrey S. Walz. 1998. "Modeling Judicial Federalism: Predictors of State Court Protections of Defendants' Rights under State Constitutions, 1969-1989." *Publius* 28:43–59.

Black, Ryan C., Ryan J. Owens, Justin Wedeking and Patrick C. Wohlfarth. 2016. *U.S. Supreme Court Opinions and Their Audiences*. Cambridge University Press.

Bonica, Adam and Michael J. Woodruff. 2014. "A Common-Space Measure of State Supreme Court Ideology." *Journal of Law, Economics & Organization* Forthcoming.

Britton, Bruce K., Shawn M. Glynn, Bonnie J.F. Meyer and M.J. Penland. 1982. "Effects of Text Structure on Use of Cognitive Capacity During Reading." *Journal of Educational Psychology* 74(1):51–61.

Caldeira, Gregory A. 1985. "The Transmission of Legal Precedent: A Study of State Supreme Courts." *American Political Science Review* 79(1):178–94.

Corley, Pamela C. 2008. "The Supreme Court and Opinion Content: The Influence of Parties' Briefs." *Political Research Quarterly* 61:468–478.

Corley, Pamela C and Justin Wedeking. 2014. "The (Dis) Advantage of Certainty: The Importance of Certainty in Language." *Law & Society Review* 48(1):35–62.

Corley, Pamela C., Paul M. Collins and Bryan Calvin. 2011. "Lower Court Influence on U.S. Supreme Court Opinion Content." *Journal of Politics* 73:31–44.

Cross, Frank B., James F. Spriggs, II, Timothy R. Johsnson and Paul J. Wahlbeck. 2010. "Citations in the U.S. Supreme Court: An Empirical Study of Their Use and Significance." *University of Illinois Law Review* 2010(2):489–576.

Epstein, Lee, Andrew D. Martin, Jeffrey A. Segal and Chad Westerland. 2007. "The Judicial Common Space." *Journal of Law, Economics, & Organization* 23(2):303–325. Data available at http://epstein.usc.edu/research/JCS.html.

Flesch, Rudolf. 1948. "A New Readibility Yardstick." *Journal of Applied Psychology* 32:221–233.

Fowler, J.H., T.R. Johnson, J.F.S. II, S. Jeon, P.J. Wahlbeck and O.B. Drive. 2007. "Network Analysis and the Law: Measuring the Legal Importance of Supreme Court Precedents." *Political Analysis* 15(3):324–346.

Friedman, Barry. 2006. "Taking Law Seriously." *Perspectives on Politics* 4(2):261–276.

Galanter, Marc. 2004. "The Vanishing Trial: An Examination of Trials and Related Matters in Federal and State Courts." *Journal of Empirical Legal Studies* 1(3):459–570.

Garner, Bryan A. 2014. "Numerical Pollution: Textual citations make legal writing onerous, for lawyers and nonlawyers alike." *ABA Journal* Feb. 1.

Gunning, Robert. 1952. *The Technique of Clear Writing.* McGraw-Hill International Book Company.

Hansford, Thomas G. and Chelsea Coe. 2014. "Linguistic Complexity and Public Acceptance of Supreme Court Decisions." Paper Presented at the 2014 Annual Meeting of the American Political Science Association. Available at `http://faculty.ucmerced.edu/thansford/Working%20Papers/Hansford_Coe_Complexity_APSA.pdf`.

Hansford, Thomas G. and James F. Spriggs, II. 2006. *The Politics of Precedent on the U.S. Supreme Court.* Princeton University Press.

Hinkle, Rachael K. 2015. "Legal Constraint in the U.S. Courts of Appeals." *Journal of Politics* 77:721–735.

Hinkle, Rachael K., Andrew D. Martin, Jonathan D. Shaub and Emerson H. Tiller. 2012. "A Positive Theory And Empirical Analysis Of Strategic Word Choice In District Court Opinions." *Journal of Legal Analysis* 4:407–444.

Knight, Jack. 2009. "Are Empiricists Asking the Right Questions about Judicial Decision-making?" *Duke Law Journal* 58:1531–1556.

Maltzman, Forrest, James F. Spriggs, II and Paul J. Wahlbeck. 2000. *Crafting Law on the Supreme Court: The Collegial Game*. Cambridge University Press.

Nelson, Michael J. 2014. "Elections and Explanations: Judicial Elections and the Readability of Judicial Opinions." Paper Presented at the 2014 Meeting of the American Political Science Association. Available at `http://mjnelson.org/papers/NelsonReadability.pdf`.

Niblett, Anthony. 2010. "Do Judges Cherry Pick Precedents to Justify Extra-Legal Decisions? A Statistical Examination." *Maryland Law Review* 70:234.

Owens, Ryan J. and Justin P. Wedeking. 2011. "Justices and Legal Clarity: Analyzing the Complexity of US Supreme Court Opinions." *Law & Society Review* 45(4):1027–1061.

Owens, Ryan J., Justin P. Wedeking and Patrick C. Wohlfarth. 2013. "How the Supreme Court Alters Opinion Language to Evade Congressional Review." *Journal of Law and Courts* 1(1):35–59.

Posner, Richard A. 2001. "Against Footnotes." *Court Review* Summer:24–25.

Resnik, Judith. 1982. "Managerial Judges." *Harvard Law Review* 96(2):374–448.

Schauer, Frederick. 2008. "Authority and Authorities." *Virginia Law Review* pp. 1931–1961.

Smith, E.A. and R.J. Senter. 1967. *Automated Readability Index*. Aerospace Medical Research Laboratories.

Spriggs, II, James F. and Thomas G. Hansford. 2000. "Measuring Legal Change: The Reliability and Validity of *Shepard's* Citations." *Political Research Quarterly* 53(2):327–341.

Wahlbeck, Paul J. 1997. "The Life of the Law: Judicial Politics and Legal Change." *Journal of Politics* 59(3):778–802.

Supporting Information

Appendix A: Measuring Readability

This appendix explains the readability measures used in our analyses. We measured the readability of the opinions using several different measures of readability, and then we combined the measures using a factor analysis. In this appendix, we first explain the formulas we used to measure opinion readability. Then, we describe the results of the factor analysis and provide some additional evidence about the validity of the measures.

The Flesch Reading Ease Scale (FRES) measures the readability of a text on a scale from 0-100 with higher scores for a text indicating greater ease of understanding (Flesch 1948).

$$\text{FRES} = 206.835 - 1.015\left(\frac{\text{Total Words}}{\text{Total Sentences}}\right) - 84.6\left(\frac{\text{Total Syllables}}{\text{Total Words}}\right) \tag{1}$$

Scores less than 30 are typically readable to individuals with a college degree, and scores ranging from 60-70 are readable by the average teenager.

The Flesch-Kincaid Grade Level (FKGL) measures the number of years of education typically required to read a text (Flesch 1948):

$$\text{FKGL} = 0.39\left(\frac{\text{Total Words}}{\text{Total Sentences}}\right) + 11.8\left(\frac{\text{Total Syllables}}{\text{Total Words}}\right) - 15.59 \tag{2}$$

The Gunning-Fox Index (Gunning 1952) uses a slightly different formula to measure the same quantity of interest as the FKGL:

$$\text{FOG} = 0.4\left[\left(\frac{\text{Total Words}}{\text{Total Sentences}}\right) + 100\left(\frac{\text{Number of Words with 3+ Syllables}}{\text{Total Words}}\right)\right] \tag{3}$$

The Simple Measure of Gobbledygook (SMOG) is another measure of the number of years

of education necessary to read a text:

$$\text{SMOG} = 1.0430\sqrt{3\text{+ Syllable Words} \times \frac{30}{\text{Total Sentences}}} + 3.1291 \tag{4}$$

The Automated Readability Index (Smith and Senter 1967) also provides a grade-level estimate of the difficulty of the text:

$$\text{ARI} = 4.71\left(\frac{\text{Total Characters}}{\text{Total Words}}\right) + 0.5\left(\frac{\text{Total Words}}{\text{Total Sentences}}\right) - 21.43 \tag{5}$$

Finally, we included a simplistic measure of readability: the average sentence length:

$$\text{ASL} = \frac{\text{Total Words}}{\text{Total Sentences}} \tag{6}$$

Having estimated the readability of each opinion using each measure, we combined the six separate indicators into a single measure of readability using a confirmatory factor analysis. The results indicate that the quantity of interest is strongly unidimensional (the eigenvalue of the second factor is only 0.34), and each of the indicators loads onto the factor at a level above 0.80. Table 3 provides the factor loadings.

Indicator	Loading
FRES	-0.94
FKGL	0.99
FOG	0.99
SMOG	0.98
ARI	0.82
ASL	0.89

Table 3: Factor Analysis Loadings. This table provides the loadings of each of the indicators of readability on the first factor extracted using confirmatory factor analysis.

Table 4 provides evidence of the validity of these indicators, along with a summary of the metric one should use to interpret each measure. Specifically, we provide the estimated readability of three classic texts: Dr. Seuss's *The Cat in the Hat*, L. Frank Baum's *The Wizard of Oz* and Lewis Carroll's *Alice in Wonderland*. All of the measures recognize *The*

Cat in the Hat is a text geared for elementary school students. Similarly, all of the measures classify *The Wizard of Oz* as a book geared toward middle school students while *Alice in Wonderland* is written for students in high school. The calculated factor scores for these three books (-58.33, -44.28, and -14.56) correspond to the first grade, sixth grade, and college freshman reading levels, respectively. This provides additional evidence of the validity of the variable.

Measure	Interpretation	*Cat in the Hat*	*Wizard of Oz*	*Alice in Wonderland*
FRES	> -30 College Degree	-106.2	-84.2	-61.2
FKGL	Grade Level	0.2	4.2	12.2
FOG	Grade Level	2.9	6.3	14.2
SMOG	Grade Level	4.0	7.3	10.5
ASL	Higher = More Complex	6.0	10.0	29.0
ARI	Grade Level	0.4	7.1	16.6

Table 4: Summary of Readability Measures with Facial Validity Checks

Appendix B: Institution-Specific Regression Results

	SCOTUS	**Same Ct.**	**Sister Cts.**	**Lower Cts.**
Readability	−0.147	0.059∗	0.047	0.048∗
	(0.084)	(0.018)	(0.030)	(0.019)
Footnotes	−0.011	−0.006∗	−0.008∗	−0.004∗
	(0.008)	(0.002)	(0.003)	(0.002)
Non-Binding Cases Cited	0.032∗	−0.011∗	0.029∗	0.006∗
	(0.006)	(0.002)	(0.003)	(0.002)
Elite Law School	−0.172	0.112∗	0.169∗	−0.032
	(0.246)	(0.042)	(0.070)	(0.036)
Subject Expertise	−0.526∗	−0.215∗	−0.494∗	−0.022
	(0.254)	(0.040)	(0.076)	(0.036)
Total Citations Received	0.000	0.004∗	0.002∗	0.009∗
	(0.000)	(0.000)	(0.000)	(0.001)
Logged Word Count	0.788∗	0.755∗	0.806∗	0.526∗
	(0.252)	(0.048)	(0.083)	(0.047)
Binding Cases Cited	−0.010	0.003	−0.014∗	−0.004∗
	(0.006)	(0.001)	(0.003)	(0.001)
Court Baseline Rate	0.005∗	0.001∗	0.005∗	0.005∗
	(0.001)	(0.001)	(0.001)	(0.001)
Author Ideological Extremity	0.004	−0.000	−0.001	−0.003∗
	(0.004)	(0.001)	(0.001)	(0.001)
Years Available	0.029	0.063∗	0.098∗	0.030∗
	(0.035)	(0.009)	(0.012)	(0.008)
Constant	−11.688∗	−6.308∗	−8.958∗	−4.001∗
	(2.110)	(0.368)	(0.645)	(0.397)
N	7503	7503	7503	7503

Table 5: Discussion by Institutional Context. Negative binomial regression estimates of the effect of opinion and author characteristics on the total number of times an opinion is discussed . . . Robust standard errors, clustered on the precedent court and year, are reported in parentheses below each coefficient, and * denotes a p-value less than 0.05.

	Positive Treatment			Negative Treatment		
	Same	**Sister**	**Lower**	**Same**	**Sister**	**Lower**
Readability	0.035	0.045	0.131*	−0.026	0.034	0.008
	(0.019)	(0.030)	(0.021)	(0.027)	(0.037)	(0.030)
Footnotes	−0.006*	−0.010*	−0.003*	−0.003	0.001	−0.006*
	(0.002)	(0.003)	(0.001)	(0.003)	(0.003)	(0.002)
Non-Binding Cases Cited	−0.014*	0.021*	−0.003	0.005	0.018*	0.018*
	(0.002)	(0.003)	(0.002)	(0.003)	(0.004)	(0.003)
Elite Law School	0.059	0.217*	−0.026	0.081	0.020	0.064
	(0.043)	(0.059)	(0.033)	(0.053)	(0.078)	(0.056)
Subject Expertise	−0.137*	−0.276*	−0.052	−0.101	−0.545*	0.021
	(0.038)	(0.062)	(0.033)	(0.056)	(0.089)	(0.056)
Total Citations Received	0.003*	0.001*	0.008*	0.000	0.000*	0.001
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
Logged Word Count	0.628*	0.498*	0.462*	0.358*	0.350*	0.264*
	(0.040)	(0.072)	(0.042)	(0.064)	(0.089)	(0.079)
Binding Cases Cited	0.003*	−0.003	−0.004*	−0.002	−0.007*	−0.003
	(0.001)	(0.002)	(0.001)	(0.002)	(0.003)	(0.002)
Court Baseline Rate	0.000	0.002*	0.004*	0.001	0.004*	0.004*
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Author Ideological Extremity	0.000	−0.000	−0.002*	−0.000	−0.002	−0.004*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Years Available	0.033*	0.082*	−0.020*	0.105*	0.111*	0.047*
	(0.009)	(0.009)	(0.007)	(0.012)	(0.013)	(0.011)
Other Type of Treatments	0.233*	0.577*	0.122*	0.099*	0.424*	0.026*
	(0.018)	(0.044)	(0.011)	(0.008)	(0.035)	(0.003)
Constant	−5.168*	−6.458*	−2.926*	−5.230*	−6.282*	−3.548*
	(0.325)	(0.575)	(0.342)	(0.500)	(0.697)	(0.660)
N	7503	7503	7503	7503	7503	7503

Table 6: Treatment by Institutional Context. Negative binomial regression estimates of the effect of opinion and author characteristics on the total number of times an opinion is discussed . . . Robust standard errors, clustered on the precedent court and year, are reported in parentheses below each coefficient, and * denotes a p-value less than 0.05.