University of Chicago Law School
Chicago Unbound

Public Law and Legal Theory Working Papers | Working Papers

2011

Dustin Guzior and Anup Malani, Implicit Bias in Legal Interpretation

Ward Farnsworth

Follow this and additional works at: http://chicagounbound.uchicago.edu/public_law_and_legal_theory

Part of the Law Commons

Recommended Citation
Ward Farnsworth, "Dustin Guzior and Anup Malani, Implicit Bias in Legal Interpretation" (University of Chicago Public Law & Legal Theory Working Paper No. 365, 2011).

This Working Paper is brought to you for free and open access by the Working Papers at Chicago Unbound. It has been accepted for inclusion in Public Law and Legal Theory Working Papers by an authorized administrator of Chicago Unbound. For more information, please contact unbound@law.uchicago.edu.

University of Chicago Institute for Law
Law & Economics, Olin Paper No. 577
Public Law & Legal Theory Research Paper No. 365

~and~

Boston University School of Law
Law & Economics Research Paper No. 11-54
Public Law & Legal Theory Research Paper No. 11-54

IMPLICIT BIAS IN LEGAL INTERPRETATION

Ward Farnsworth, Dustin Guzior and Anup Malani

This paper can be downloaded without charge from:
John M. Olin Program in Law and Economics Working Paper Series:
http://www.law.uchicago.edu/Lawecon/index.html
Public Law and Legal Theory Working Paper Series:
http://www.law.uchicago.edu/academics/publiclaw/index.html

or from the Boston University site at:
http://www.bu.edu/law/faculty/scholarship/workingpapers/2011.html;
http://www.bu.edu/law/faculty/scholarship/workingpapers/Farnsworth-Guzior-Malani.html;
http://www.bu.edu/law/faculty/scholarship/workingpapers/documents/FarnsworthW-GuziorD-MalaniA102611.pdf.

and The Social Science Research Network Electronic Paper Collection.

IMPLICIT BIAS IN LEGAL INTERPRETATION

Ward Farnsworth, Dustin Guzior and Anup Malani[1]

What role do policy preferences play when a judge or any other reader decides what a statute or other legal text means? Most judges think of themselves as doing law, not politics. Yet the observable decisions that judges make often follow patterns that are hard to explain by anything other than policy preferences. Indeed, if one presses the implications of the data too hard, it is likely to be heard as an accusation of bad faith—a claim that the judge or other decision-maker isn't really earnest in trying to separate preference from judgment. This does not advance the discussion, and distracts from the possibility of more interesting explanations. A promising antidote, we believe, lies in empirical study not just of large numbers of judicial decisions collected over time, as previous scholars have done, but of the immediate experience of legal interpretation.

We compile, and here present, rich evidence of what happens when lawyers in training are asked in controlled surveys to distinguish between their policy preferences on the one hand and their own interpretive judgments or predictions about courts judgments on the other. Our findings offer two lessons. First and foremost, they suggest that separating policy preferences from judgments about the meaning of statutes is very difficult. The same is true of preferences and predictions about what courts will do: respondents tend to predict that courts will do what the respondents themselves prefer. The fundamental entanglement of preferences and interpretation raises important questions about the ability of anyone – including judges – to neutrally carry out interpretive strategies meant to generate answers in close cases. Second, however, the results also show that certain ways of framing the interpretive question can reduce the influence of preference on interpretation—though perhaps not its effect on predictions. Instead of simply asking respondents how they would interpret the text of a statute or how the drafters would likely want it applied, it is better to ask respondents how ordinary readers would interpret the statute. This framing of the interpretative question can debias an individual's interpretation of a statute.

In short, interpretative theories that elevate text alone or give the intent of drafters are both susceptible to contamination by private preferences. To immunize interpretation from these preferences, a theory that asks how ordinary readers would read a statute may be the best prescription.

I. Methodology 4
A. Cases 5
B. Questions 7
II. Results 8
A. The influence of preferences on interpretation 8
1. Ordinary meaning 11
2. Drafters' intent 12
3. Ordinary readers 13
B. The influence of preferences on predictions 15
III. Implications 18

---

[1] Farnsworth is Associate Dean for Academic Affairs and Professor of Law at Boston University. Guzior is an associate at Sullivan & Cromwell. Malani is Professor Law at the University of Chicago, and affiliated with Resources for the Future, and the National Bureau of Economic Research. We thank Frank Easterbrook, Einer Elhauge, William Eskridge, Richard Posner, Stephen Williams, seminar participants at the University of Chicago for helpful comments. We thank Adam Badawi, Anthony Casey, Mary Anne Franks, Adam Muchmore, Anthony Niblett, and Arden Rowell for helping us administer surveys at the University of Chicago Law School.

Electronic copy available at: http://ssrn.com/abstract=1944273

IV. Limitations of the study 20
A. Causation 20
B. Criminal cases 20
C. External validity 21
Conclusion 24

INTRODUCTION

What role do policy preferences play when a judge or any other reader decides what a statute or other legal text means? This is a stubborn problem in the study of legal interpretation, and it continues to provoke blunt disagreement at the highest levels of the profession. The Chief Justice of the United States says that his job when deciding cases is to "call balls and strikes";[2] his claim is ridiculed by distinguished legal scholars and others.[3] Prominent academics publish studies suggesting that judges' views of policy play an important role in their decisions;[4] prominent judges deny it and complain that the studies do not take seriously enough the reports that the judges themselves make of their own thinking.[5]

Even within the academy, those studies suggesting that close cases are decided according to the ideologies of the judges are open to controversy. Academics who accept the influence of ideology debate whether its effect is large or small,[6] and meanwhile much scholarship on legal interpretation continues to be premised on the opposite vision.[7] That scholarship considers at length when only the text of a statute should be considered by a judge, and when other evidence should be taken into account—all of this based impliedly or explicitly on the idea that judges might carry out such decisions in a straightforward way that does not give effect to their preferences.

---

[2] Confirmation Hearing on the Nomination of John G. Roberts, Jr. to Be Chief Justice of the United States: Hearing Before the S. Comm. on the Judiciary, 109th Cong. 56 (2005).

[3] See Richard A. Posner, *Some Realism About Judges: A Reply to Edwards and Livermore*, 59 Duke L.J. 1177 (2010) (calling Roberts's analogy "ludicrous" and "ridiculous"); Erwin Chemerinsky, Seeing the Emperor's Clothes: Recognizing the Reality of Constitutional Decision Making, 86 B.U. L. Rev. 1069, 1069 (2006) ("it is hard to think of a less apt analogy").

[4] See, e.g., Richard L. Revesz, *Environmental Regulation, Ideology, and the D.C. Circuit*, 83 Va. L. Rev. 1717 (1997); Frank B. Cross & Emerson H. Tiller, *Judicial Partisanship and Obedience to Legal Doctrine: Whistleblowing on the Federal Courts of Appeals*, 107 Yale L.J. 2155 (1998); Cass R. Sunstein et al., Are Judges Political?: An Empirical Analysis of the Federal Judiciary 150 (2006).

[5] Harry T. Edwards, *Collegiality and Decision Making on the D.C. Circuit*, 84 Va. L. Rev. 1335 (1998); Patricia M. Wald, *A Response to Tiller and Cross*, 99 Colum. L. Rev. 235 (1999); Harry T. Edwards and Michael A. Livermore, *Pitfalls of Empirical Studies That Attempt to Understand the Factors Affecting Appellate Decisionmaking*, 58 Duke L.J. 1895 (2009).

[6] See, e.g., Gregory C. Sisk and Michael Heise, *Judges and Ideology: Public and Academic Debates About Statistical Measures*, 99 Nw. U. L. Rev. 743 (2005); Frank B. Cross, Decision Making in U.S. Courts of Appeals (2007).

[7] See, e.g., Antonin Scalia, A Matter of Interpretation: Federal Courts and the Law (1997); Jonathan T. Molot, *The Rise and Fall of Textualism*, 106 Colum. L. Rev. 1 (2006); John F. Manning, *What Divides Textualists from Purposivists?* 106 Colum. L. Rev. 70 (2006).

Electronic copy available at: http://ssrn.com/abstract=1944273

These conflicts about how judges do their job are stubborn for understandable reasons. Most judges think of themselves as doing law, not politics. They do not want to give effect to their policy preferences, so they decide not to, and they think that they aren't. That inner feeling of legalism is convincing and tenacious, and it is not limited to judges. It is common enough for most people deciding contentious legal questions to feel that they are doing law and that their adversaries are doing politics. Yet the resulting decisions that judges make often follow patterns that are observable from the outside and seem hard to explain by anything other than policy preferences. There is, in other words, a discrepancy between the subjective experience of legal decision-making and objective observations of the results. Subjective feelings are slow to yield in such a contest, so it is no surprise that judges dispute the implications of objective studies. Indeed, if one presses the implications of the data too hard, it is likely to be heard as an accusation of bad faith—a claim that the judge or other decision-maker isn't really earnest in trying to separate preference from judgment. This does not advance the discussion, and distracts from the possibility of more interesting explanations.

A promising antidote, we believe, lies in empirical study not just of large numbers of judicial decisions collected over time, as previous scholars have done, but of the immediate experience of legal interpretation. What happens when people are asked to interpret legal texts while keeping their preferences out of it? If we take simple tasks of separation—here, separation of preference from judgment in a legal setting—and we instruct people to carry them out, how well can they do it? Do their preferences affect their readings anyway?

These are not specifically questions about what judges can do. They are questions about what anyone can do; they are questions about cognition, not the judicial role. But they may have implications for the judicial role. Judges are human, and what they (or anyone) *can* do has implications for the scope of their charge and how they should try to carry it out. Moreover, the questions about bias and cognition are central to—even if they remain largely latent in—academic debates over legal interpretation. Mostly they bubble just under the surface and as yet they have been the subject of no direct inquiry.

This Article starts that project. It cannot directly settle the question of how judges carry out those tasks, because large-scale controlled experiments involving judges are difficult to carry out. We instead sought to compile, and here present, rich evidence of what happens when lawyers in training are asked in controlled surveys to distinguish between their policy preferences on the one hand and their own interpretive judgments or predictions about courts judgments on the other.

Our findings offer two lessons. First and foremost, they suggest that separating policy preferences from judgments about the meaning of statutes is very difficult. The same is true of preferences and predictions about what courts will do: respondents tend to predict that courts will do what the respondents themselves prefer. The fundamental entanglement of preferences and interpretation raises

important questions about the ability of judges (or anyone) to neutrally carry out interpretive strategies meant to generate answers in close cases. It also illustrates subtle and significant challenges to the lawyer's role in attempting to provide accurate advice about the law to clients.

Second, however, the results also show that certain ways of framing the interpretive question can reduce the influence of preference on interpretation—though perhaps not its effect on predictions. The key is to start with a non-idealized, external reference point for making the interpretive judgment. Instead of simply asking respondents how they would interpret the text of a statute or how the drafters would likely want it applied, respondents are better asked how ordinary readers would interpret statute. When this is done, the influence of a respondent's preferences on interpretation subsides. In short, proper framing of the interpretative question can debias an individual's interpretation of a statute.

Our findings, then, are not only cautionary; they also have suggestions to offer about how to reduce unwanted influences on anyone's reasoning about what an ambiguous statute means. Interpretative theories that elevate text alone or give the intent of drafters are both susceptible to contamination by private preferences. To immunize interpretation from these preferences, a theory that asks how ordinary readers would read a statute may be the best prescription. Unfortunately there is no comparably straightforward remedy for bias in predictions about the behavior of judges. And since common models of settlement suggest that conflicting predictions among parties drive litigation behavior, our findings suggest a new explanation for why parties often fail to settle.

Section I of this paper describes our survey instrument. Section II describes the results. Section III discusses the implications for interpretative theory. Section IV addresses some limitations, and the final section offers concluding thoughts.

I. METHODOLOGY

We proceeded by creating survey instruments and administering them to over 1500 law students, most of them in their first semester of study.[8] In each survey, the respondent was presented with ambiguous statutes and facts to which they might apply. The statutes and facts were taken from Supreme Court cases involving federal criminal law or civil disputes. The respondents were told what position each side to the case took.

The questions then put to them took various forms. First, the respondents were asked what outcome of the case they preferred as a matter of policy, setting aside the text of the statute. Then they were asked to interpret the text of the statute, setting aside their policy preferences. Each respondent was asked this question in

[8] The latest surveys were administered to first year law students at Boston University and University of Chicago during the 2010-2011 academic year.

different ways. Some were asked which reading of the text they thought was the best fit to its ordinary meaning; others were asked which reading ordinary readers of English would think the best fit to its ordinary meaning; and yet others were asked which reading they thought was the best fit to the drafters' intent. Finally, respondents were asked on occasion to predict which reading a court would prefer. The goal, of course, was to find any relationships between answers to the question about policy preferences and answers to the interpretive and predictive questions. As we shall see, policy preferences seem to affect them all—but not to the same extent.

A. *Cases*

Each of our surveys contained questions about a number of different statutory cases. The results from the surveys were highly repetitive from case to case, so we review in detail four of the cases from our most recent survey instruments.

The first case ("Gun use case") was based on *Smith v. United States*, 508 U.S. 223 (1993).[9]

> A federal statute, 18 U.S.C. § 924(c), provides an enhanced prison sentence for anyone who "uses" a firearm "in relation to . . . a drug trafficking offense." Defendant, a drug dealer, owned a gun. He approached a confederate and offered to trade him the gun for some cocaine. His confederate turned out to be an undercover police officer, and defendant was arrested. He was charged with violating 924(c). Defendant did not brandish the gun or use it in a threatening manner, but he did offer it as an item of barter.
>
> The question is whether offering the gun in trade was a "use" of it within the meaning of 924(c) (in which case the defendant gets the extra time on his prison sentence). Defendant's reading is that offering a gun in trade is not a "use." The government's reading is that it is a "use."

The second case ("Child pornography case"), adapted from *United States v. X-Citement Video, Inc.*, 513 U.S. 64 (1994),[10] was somewhat more complicated.

> A federal statute, 18 U.S.C. § 2252, reads in part as follows:
>
> (a) Any person who—
>
> (1) knowingly transports or ships in interstate or foreign commerce by any means including by computer or mails, any visual depiction, if—

[9] The Court held, 6-3, in favor of the government.

[10] The Court held, 7-2, that the government had to prove that the defendant knew that the films he sold included sexually explicit acts by minors. We presented the case to our respondents in a form a bit different, and a bit simpler, than the form it took in the Supreme Court. In the actual *X-Citement Video* case, it was the defendant who argued that the scienter requirement did not reach the age of the performers in the movies—because he claimed this made the statute unconstitutional. Since we did not wish to engage the constitutional question, we wrote the survey question to suggest that the defendant argued for a reading of the statute that made it harder to get a conviction under it.

> (A) the producing of such visual depiction involves the use of a minor engaging in sexually explicit conduct; and
>
> (B) such visual depiction is of such conduct;
>
> shall be punished as provided in subsection (b) of this section.

The defendant was accused of violating the statute by selling pornographic videotapes that included footage of a woman who was under the age of 18, and thus was a minor. He defended on the ground that when he sold the tape, he did not know the person on the tape was a minor.

The question is whether the word "knowingly" in section (1) applies to the phrase "the use of a minor" in section (1)(a). The defendant's reading is that "knowingly" does modify "the use of a minor." The government's reading is that "knowingly" does not modify "the use of a minor."

The third case ("False statements case") was based on *United States v. Yermian*, 468 U.S. 63 (1984).[11]

The federal "false statements" statute, 18 U.S.C. § 1001, says:

> Whoever knowingly and willfully falsifies, conceals or covers up by any trick, scheme, or device a material fact, or makes any false, fictitious or fraudulent statements or representations, or makes or uses any false writing or document knowing the same to contain any false, fictitious or fraudulent statement or entry, if the matter lies within the jurisdiction of any department or agency of the United States, shall be fined not more than $10,000 or imprisoned not more than five years, or both.

The defendant worked for a company that had a contract with the Department of Defense. The company asked him to fill out a questionnaire to obtain a security clearance. He did so. His company mailed the questionnaire to the Department of Defense. The Department discovered that the defendant's answers contained false statements. He was charged with violating the statute quoted above. His defense was that he had not realized that his questionnaire would be forwarded to the government.

The question is whether the statute requires proof that a defendant knew the matter in question was within the jurisdiction of a government agency. The defendant's reading is that the statute does require such proof. The government's reading is that it does not require such proof.

The final case ("Attorney's fees" case) was a civil case based on *West Virginia Univ. Hospitals v. Casey*, 499 U.S. 83 (1991).

A federal statute, 42 U.S.C. § 1988, provides in relevant part: "In any action or proceeding to enforce a provision of section 1983 of this title, the court, in its discretion, may allow the prevailing party [to recover from the losing party] a reasonable attorney's fee as part of its costs." Plaintiff brought a successful lawsuit against the government to enforce section 1983 and sought to recover fees paid to experts who advised his attorney.

---

[11] The Court held for the government, 5-4, that knowledge of the federal agency's jurisdiction on Yermian's part was not needed to support his conviction.

> The question is whether fees paid to experts by an attorney are covered by the part of § 1988 allowing recovery of "a reasonable attorney's fee." The defendant's reading is that fees of experts who advise an attorney are not covered by § 1988. The plaintiff's reading is that experts' fees are covered.

B. *Questions*

After presenting facts from one of the cases above, the survey asked each respondent about her policy preference:

> Setting aside the text of the statute, who do you think should win as a matter of policy preference?
>
> (A) I strongly prefer that the defendant win
> (B) I mildly prefer that the defendant win
> (C) I mildly prefer that the government win
> (D) I strongly prefer that the government win

Each respondent was also asked whether the defendant or the government's reading of the statute was better.[12] This question was asked in different ways. Some respondents were asked:

> Setting your policy preference aside, which reading better fits the ordinary meaning of the statute's text?
>
> (A) The defendant's reading.
> (B) Probably the defendant's reading.
> (C) Probably the government's reading.
> (D) The government's reading.

We will refer to the question just shown above as the "ordinary meaning" question. Some respondents were instead asked a different question that we will call "drafters' intent":

> Setting your policy preference aside, which reading of the statute is a better fit to what the drafters of the statute intended?
>
> (A) The defendant's reading.
> (B) Probably the defendant's reading.
> (C) Probably the government's reading.
> (D) The government's reading.

Finally other respondents were asked an "ordinary readers" question:

> Setting your policy preference aside, which reading of the statute would ordinary readers of English think is a better fit to the ordinary meaning of the statute's text?

---

[12] In describing the survey we present the interpretive question after the preference question. Later we will present the predictive question. When we administered the survey, however, we mixed up the order of the questions across cases. Thus, for some cases the preference question followed the interpretive question and for others the predictive question preceded the other two. We did this to determine if the order of questions affected answers. We found that it did not.

(A) The defendant's reading.
(B) Probably the defendant's reading.
(C) Probably the government's reading.
(D) The government's reading.

Finally, after the gun use case in particular, respondents were asked about how they predict a court would interpret the statute:

Which side's reading do you predict that a court would agree with?

(A) The defendant's reading.
(B) Probably the defendant's reading.
(C) Probably the government's reading.
(D) The government's reading.

We thus recorded policy preferences for everyone, and examined the relationship between those preferences and the answers given to the other questions.

II. RESULTS

A. *The influence of preferences on interpretation*

Figures 1-4 describe for the four cases the correlation between respondents' preferences and their answers to the different versions of the interpretive question. The horizontal axis lines up respondents according to their policy preferences about the outcome of the case (from those who strongly preferred that the defendant win to those who strongly preferred that the government win). The vertical axis shows which side's reading the respondents thought were best in reply to the various interpretive questions we asked. We coded the reading from 1 to 4, where 4 indicates the most the most pro-government. Using this scale, each line presents the average answer to an interpretive question among respondents with a given policy preference. The whiskers present the 95% confidence interval for each average.

The basic pattern is obvious enough. Respondents' judgments about the ordinary meaning of the statute's text and about the drafters' intent are highly correlated with their policy preferences—even though the respondents were instructed to set those preferences aside. The one remarkable exception is the question about which side's interpretation *ordinary readers* would think better fits the ordinary meaning of the statute's text. The answers to that question are significantly less correlated with respondents' policy preferences in every case. These results are confirmed in the regression analysis reported in Table 1.[13] Let us consider each of these findings in detail.

[13] The regression is described in the note under the table. The important finding is that the coefficients on the interactions between policy preference on the one hand and ordinary readers or drafters intent questions on the other hand are positive and statistically significant in each regression.

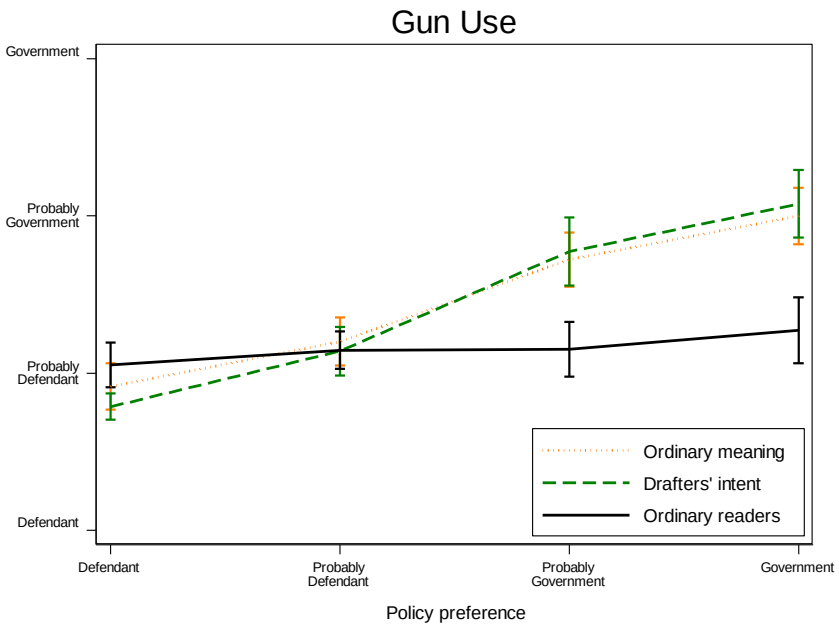

Figure 1. Correlation between policy preference and interpretation in the gun use case.

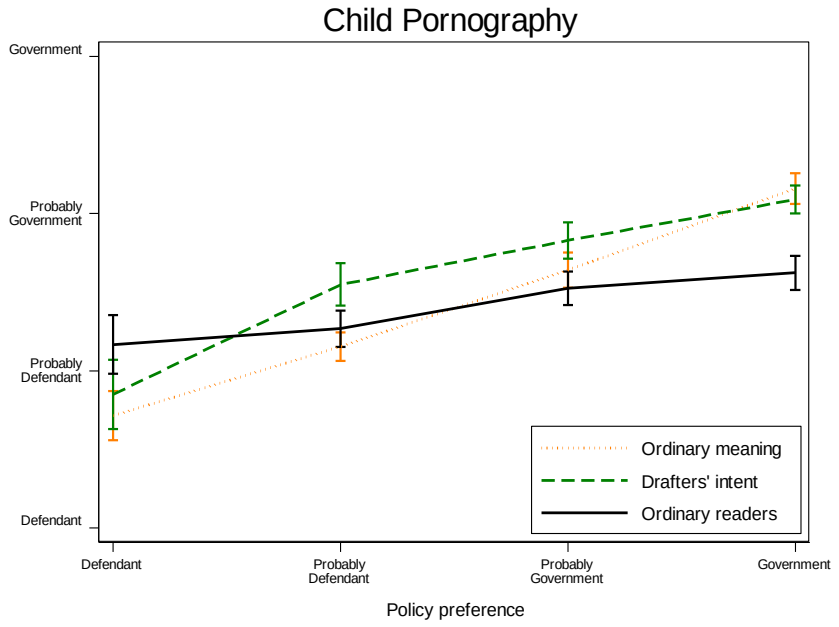

Figure 2. Correlation between policy preference and interpretation

---

This implies that in every case policy preferences have significantly more influence in the ordinary meaning and drafters' intent questions than in the ordinary readers question.

in the child pornography case.

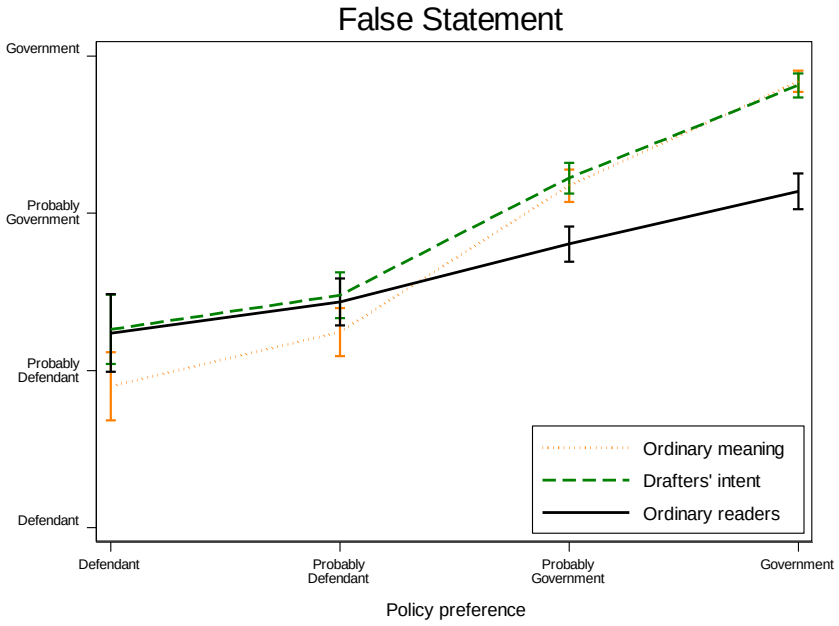

Figure 3. Correlation between policy preference and interpretation in the false statement case.

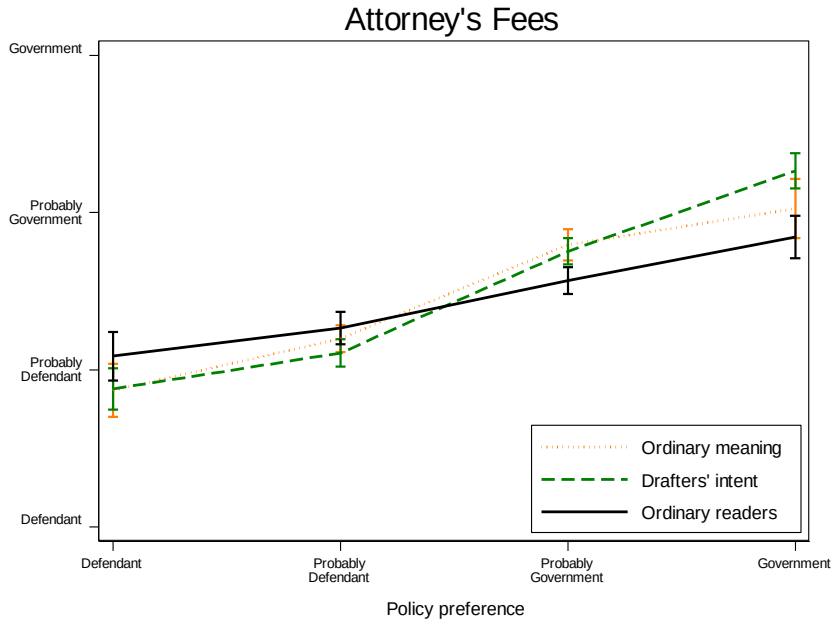

Figure 4. Correlation between policy preference and interpretation in the attorney's fees case.

Table 1. Regression analysis of correlation between policy preference and interpretation.

	Child porn	False state	Attorney fees	Gun use	All
Constant	1.975***	1.857***	1.782***	1.990***	1.840***
	(0.184)	(0.238)	(0.153)	(0.185)	(0.114)
Ordinary meaning	-0.733***	-0.774**	-0.383*	-0.477*	-0.650***
	(0.263)	(0.326)	(0.221)	(0.265)	(0.159)
Drafters' intent	-0.211	-0.262	-0.540**	-0.647***	-0.418***
	(0.269)	(0.329)	(0.221)	(0.232)	(0.151)
Policy preference	0.167***	0.316***	0.261***	0.066	0.209***
	(0.060)	(0.076)	(0.057)	(0.077)	(0.041)
Policy preference	0.308***	0.347***	0.170**	0.314***	0.322***
x Ordinary meaning	(0.085)	(0.105)	(0.083)	(0.106)	(0.057)
Policy preference	0.176**	0.195*	0.236***	0.373***	0.269***
x Drafters' intent	(0.087)	(0.108)	(0.081)	(0.103)	(0.054)
Number of observations	796	446	796	450	1,786
Adjusted R2	0.121	0.249	0.152	0.128	0.184

Note. Table reports results from a regression of answers to the interpretive question on indicators for how the interpretive question was framed, policy preference, and indicators for framing interacted with policy preference. Omitted category is ordinary readers question. Interpretation and policy preference are measured on a 1 - 4 scale, with 4 most favoring the government. The regression was run once for answers for each case (first 4 columns), and then for answers from all cases pooled together (last column). Robust standard errors are reported in parentheses below coefficients. ***/**/* indicate $p < 0.01/0.05/<0.1$ without adjustment for multiple testing.

1. Ordinary meaning

The results with respect to the ordinary meaning question are stark. Judgments about the ordinary meaning of a text are highly entwined with policy preferences about the outcome of the case the text is being used to decide. What makes the finding especially striking is that the respondents were explicitly told to separate the two considerations. They could not do it.

This failure has several implications. First and most simply, personal statements about what the "ordinary meaning" of a text seems to be are highly prone to bias by the policy preferences of whoever is making the claim. Second, the makers of such claims are not likely to subjectively experience themselves as biased in this way. Their intentions were otherwise. The influence exerts itself invisibly.

2. Drafters' intent

The same pattern appears in answers to the drafters' intent question.[14] We thought it possible that separating oneself a bit from the question of the statute's meaning—being asked not what *you* think it means, but you believe the author *wanted*—might reduce the influence of one's own policy preferences. Unfortunately it does not. Estimates of the drafters' intent, like judgments about ordinary meaning, closely track the respondents' own wishes.

One reason for this result might be that respondents project their own preferences onto the legislators who they imagine drafted the bill. Here as with the first question, the projection evidently is unconscious, for again the respondents were told to put their preferences aside when answering the question.

We have wondered whether the answer to the drafters' intent question might help explain the answers to the previous one about which reading of the text best fits its ordinary meaning. Maybe one reason policy preferences are bound up with replies to the ordinary meaning question is that people try to determine ordinary meaning by guessing at what the drafters of the statute must have meant, and they can only make headway on that question by asking what they themselves would have wanted if they had been the drafters. This may be part of the story, but it can't be all of it, because the responses to the "ordinary meaning" question and the "drafters' intent" sometimes were different in significant ways.

The difference between the "ordinary meaning" question and the "drafters' intent" question is nicely illustrated in Figures 2 and 3. Looking at the left half of each chart—that is, to those respondents who preferred that the defendant win—we find that people answering the "drafters' intent" question side with the government more quickly than people just saying which view of the text better fits its ordinary meaning. In the child pornography case (Fig. 2), 89% of those respondents who strongly preferred that the defendant win thought the defendant's reading better fit the statute's "ordinary meaning." By contrast, only 67% of that group thought the defendant's reading better fit the "drafters' intent". That is a significant shift in the government's favor. Respondents who mildly preferred that the defendant win exhibited the same shift: 79% of them thought the defendant's reading better fit the statute's ordinary meaning, while only 42% of them thought the defendant's reading better fit the drafters' intent.[15] Overall, 22% of respondents who strongly preferred

---

[14] It might seem that our respondents were not in a good position to comment on the intent of those who drafted the statutes they read. They did not have any statements from the legislative history, or any information about what events caused the statute to be drafted, or any knowledge of the rest of the surrounding legal context. Still, some courts like to say that the best evidence of a legislature's intent is the words they chose to use, see, e.g., *U.S.* v. *Husted*, 545 F.3d 1240, 1246 (10th Cir. 2008), so perhaps our survey-takers were not entirely disarmed. At any rate, we meant the question mostly as a heuristic.

[15] We see the same shifts when we look at the False Statement Case (Figure 3). In that case, the shift for respondents who strongly preferred that the defendant win is from 70% ("ordinary meaning") to

that the defendant win and 37% of respondents who mildly preferred that the defendant win changed sides in their judgments about which side had the better reading when they were asked about the drafters' intent rather than about ordinary meaning. There is no similar change in position for respondents on the right half of the graph—those respondents who preferred that the government win.

The result is easy enough to state. The "drafters' intent" question makes pro-defendant respondents—but not pro-government respondents—more likely to draw conclusions contrary to their policy preferences. The reason for that result is not so clear. Perhaps there is a tendency to imagine that legislators are more aggressive than oneself in wanting to put people in jail—that if legislators were asked which reading they preferred, they would err on the side of finding a violation when conduct arguably runs afoul of the statute. Whatever the explanation, it is interesting to see evidence that judgments about ordinary meaning, at least when viewed in large sets, aren't quite the same as judgments about the drafters' intent. These evidently are experienced as related but different questions.

3. Ordinary readers

One of the most striking findings of this study is that policy preferences, as pervasive as they are, do not infect all interpretive judgments equally, and often seem to have little or no effect on answers to one question in particular: which reading an *ordinary reader* would think best fits the ordinary meaning of the statute. We might call that the objective form of the question about ordinary meaning, as opposed to the subjective earlier version that asked the respondents for their own opinion about it. When asked for a judgment about what ordinary readers would think, respondents agreed a remarkably large share of the time. This question produces the black lines in the charts above that sometimes are flat or nearly so, and that always have a lesser slope than the other lines—showing in either case a much reduced entanglement with policy preferences.

The child pornography case—the second case shown earlier—is the strongest illustration of what effect the "ordinary readers" question can have. When respondents were asked to put aside their policy preferences and say which reading they thought best fit the ordinary meaning of the statute, 85% of those who strongly preferred that the government win as a matter of policy also said that the government's reading was better (or was "probably better"). But when respondents were asked which reading *ordinary readers* would think a better fit to the text, only 37% of those who strongly wanted the government to win chosen the government's reading. On the other side of the spectrum, of those who strongly preferred that the defendant win as a matter of policy, only 11% preferred the government's reading (or "probably" preferred it) on their own account. But when asked which interpretation ordinary readers would likely think correct, the number favoring the

---

56% ("drafters' intent"), a change of 14%. The shift for respondents who mildly preferred that the defendant win is from 65% ("ordinary meaning") to 48% ("drafters' intent"), a change of 17%.

government's reading rose to 40%. A final way to see the point compares the range between the answers that different responsdents gave to different questions. Those considering the child pornography problem were asked the "ordinary readers" question, all groups of respondents, regardless of their policy preferences, favored the defendant's reading from 57%-63% of the time. When simply asked to judge for themselves which reading is better, the numbers choosing the defendant's reading ranged from 16%-89%, depending on the policy preference they reported. In short, asking what ordinary readers would think the statute meant made the respondents much more likely to give an answer that went against their rooting interests, whatever they were.

We should add two caveats. First, the "ordinary readers" question does not produce such a strong two-way shift in every case. In the gun use case, for example, most of the shifting is one-way. People who favor the government as a matter of policy are likely to shift to a judgment in favor of the defendant's reading, but there is only a little movement the other way. Most of those who prefer the defendant as a matter of policy continue to say that the defendant's reading is what ordinary readers would think correct. Nevertheless, all respondents do come to general agreement in their answers, despite continued disagreement on the policy question. The gun use case seems to be an unusual one where the question about what ordinary readers would think produces an especially high level of agreement that the defendant is right, and so calls for no movement by those who are rooting for defendant on policy grounds. The three other cases considered in this Article all produce a two-way shift when the "ordinary readers" question is asked.

Second, asking what ordinary readers would think is not a cure-all for the influence of policy preference on interpretive judgments. In some cases that question does seem to wipe out the influence entirely, but in others it merely reduces the influence of preferences by comparison to its influence on other questions. Notice, for example, that in the false statements and attorney's fees cases, the "ordinary readers" line has a rather steep slope. Indeed, it is steeper than the slope of the line produced in the gun use case when respondents there are asked about the ordinary meaning of the text. In other words, asking about "ordinary readers" in the one case is worse (from the standpoint of contamination by policy preference) than asking *any* question in the other case. But that just shows that some cases produce policy preferences that are especially hard to contain. The fact remains that in any *given* case, asking what ordinary readers would think the text means always does a better job than any other question yet found of producing an answer that is independent of policy judgments.

Why does asking what ordinary readers would think do more to filter out bias than questions about the drafters' intent or the likely views of a judge? It may be that thinking about what ordinary readers would say directs one's attention to an external benchmark—the purely conventional meaning of the words—and that the attention is thus distracted from its concerns about outcomes. It may also be the case that answers become biased when the questions have any sort of aspirational

quality. The "ordinary meaning" question asked which reading the respondent thought was *better*. The question about drafters' intent invited the respondents to think about what someone else would have wanted—but the someone else wasn't just anyone. It was a legislator, a faceless but easily idealized author of the text who the respondents might easily imagine has about the same way of looking at things than they do. And the same could be said of the questions that asked what a court would likely do. This time the respondent is asked to imagine how a judge would read the language, and again the judge is a generic but idealized figure onto whom good sense—that is, policy preferences—can readily be projected.

When they are asked what ordinary readers of English would think the text means, it may be that something like the opposite movement occurs. The respondents are asked to imagine what would be thought by a population a little duller than they are: mere ordinary readers, not the better-than-average readers that most people likely feel themselves to be. (It would be surprising if any respondents thought of themselves as *worse* than ordinary.) So when they think about how ordinary readers would interpret the law, the respondents are looking due sideways or slightly down. "*Ordinary* readers? Well, I suppose *they* would just think X." The inner experience, on this speculation, is that the reading is being "dumbed down" a little when one wonders what an ordinary reader would think. But the dumbing down, if that's what it is, is useful in an unexpected way, because it strains out a lot of the wishful thinking that spoils mental inquiries made with a more upward-looking angle.

These results are consistent with earlier work in which we examined judgments about whether a text is ambiguous.[16] In that study we found that respondents with strong policy preferences about a case were much less likely to find the statute at issue to be ambiguous—assuming they were simply asked for their opinion on that question. But when they were asked whether an ordinary reader of English would likely find the statute ambiguous, their judgments came loose from their preferences in much like the way we see here.[17] The difference is that in the prior study we were talking just about threshold judgments of whether a statute fairly admits of two interpretations—an important question in statutory cases, but still separate from the final and substantive question of what the statute means. In this study we have extended the inquiry to that substantive question of statutory meaning, and we find the impact of asking an "external" or "objective" question even more profound here than it was with respect to ambiguity.

B. *The influence of preferences on predictions*

After the gun use and attorneys' fees cases, we asked respondents to predict which reading a court would prefer. As Figures 4 and 5 illustrate, their answers

---

[16] See Ward Farnsworth, Dustin Guzior, and Anup Malani, Ambiguity About Ambiguity: An Empirical Inquiry into Legal Interpretation, 2 J. Leg. Anal. 1 (2010).
[17] Id. at XX.

tracked their policy preferences. The results were confirmed by regression analysis, as reported in Table 2. They were also the same in one other case that appears in the appendix. Again, we had speculated that asking what someone else—a judge—would think about the statute might help the respondents give an answer that was independent of their own preferences. It did not. In a way parallel to what we suggested when considering drafters' intent, it may be that respondents project their preferences onto judges when they imagine them interpreting a text.

These results may shed a bit of light on why some lawsuits fail to settle. A typical settlement is based on overlapping predictions the two sides make about how a court will decide a case. To the extent those predictions are bound up with preferences about the outcome, they are likely to diverge and shrink the bargaining range between the parties.

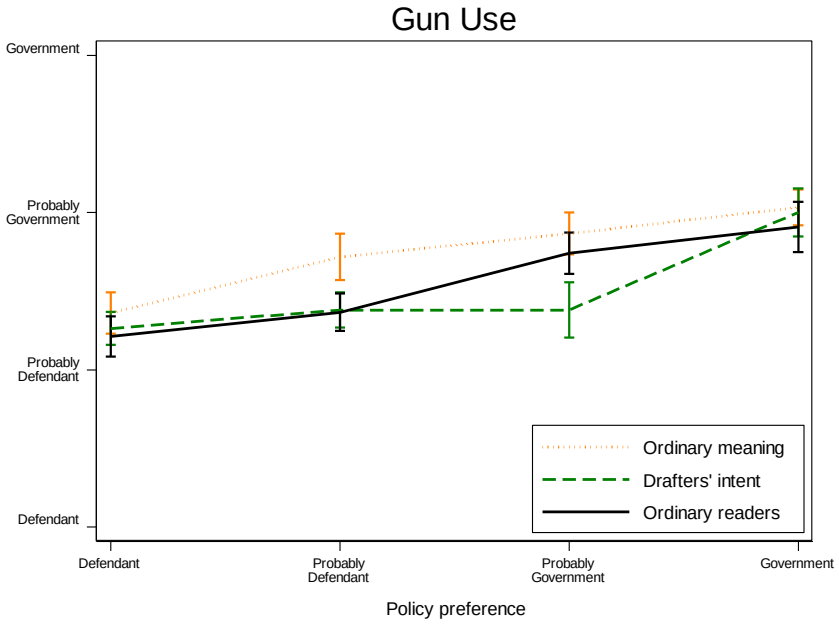

Figure 5. Correlation between policy preference and prediction in the gun use case.

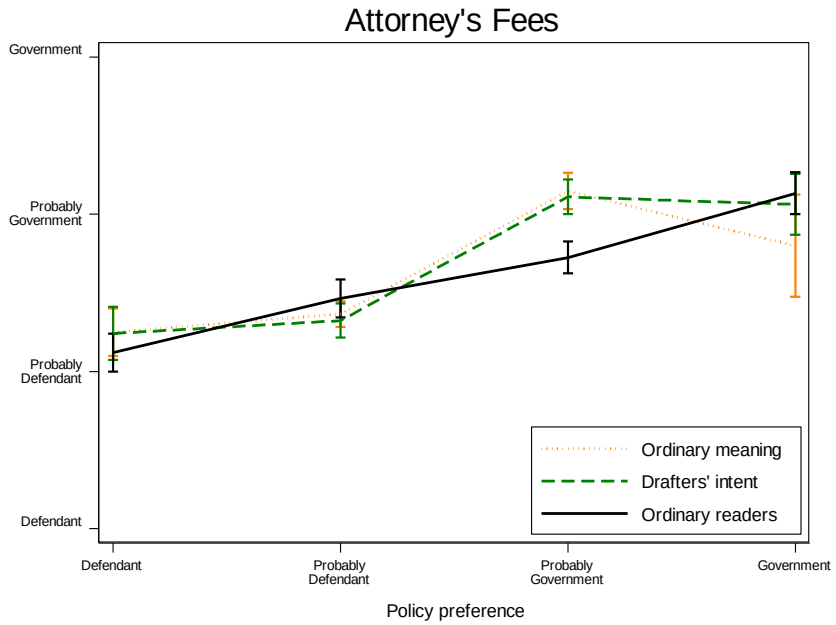

Figure 6. Correlation between policy preference and prediction in the attorneys' fees case.

Table 2. Regression analysis of correlation between policy preference and prediction.

	Attorney's Fees	Gun use
Constant	2.160***	2.044***
	(0.190)	(0.162)
Ordinary meaning	-0.029	0.222
	(0.271)	(0.231)
Drafters' intent	0.006	0.099
	(0.266)	(0.220)
Policy preference	0.157**	0.182***
	(0.072)	(0.065)
Preference x Ordinary meaning	0.016	-0.033
	(0.108)	(0.090)
Preference x Drafters' intent	0.017	-0.049
	(0.102)	(0.090)
Number of observations		505
Adjusted R2	440	0.034

Note. Regression of answers to interpretative question (1 to 4, with 4 most pro-government) on indicators for framing of interpretive question, policy preference (1 to 4, with 4 most pro-government), and preference interacted with framing indicators. Robust standard errors reported below coefficients. ***/**/* indicate p < 0.01/0.05/0.1.

III. IMPLICATIONS

Once we see the advantages of asking what ordinary readers would think a text means, further questions become apparent. First, even if answers to the "ordinary readers" question are untainted by bias, how relevant are those answers as a legal matter? Second, all worries about bias to one side, how *accurate* are the answers that respondents give when they predict what ordinary readers would think? This section addresses those two issues.

To begin with the first question, of course there are well-developed schools of thought about the goals of statutory interpretation, and it might seem possible to link some of them to choices in our survey instruments. We could suppose that when we ask which reading better comports with the drafters' intent, we are inviting the respondents to act like "intentionalists" or "purposivists."[18] And when we ask what ordinary readers of the statute would say it means, we are inviting the respondents to act like textualists—or one variety of textualist, anyway. But this picture doesn't do justice to those schools of interpretive thought. Most interpreters of statutes nowadays are likely to regard judges as agents of the legislature; they differ mostly in what evidence of legislative intent they think proper to consider.[19] Obviously a good intentionalist and a good textualist will both want a lot more material to work with than anyone received in our surveys. The respondents had no basis for comment on the legislature's purpose except their own speculations, and they didn't have all the materials bearing on "semantic context" that a textualist would want them to have.

So nothing we have found or said here strikes a great blow for any one theory of interpretation against another. But the findings are suggestive and do allow some recommendations of emphasis. Asking what an ordinary reader would think a statute means is an important part of one kind of textualism. It is the type that puts an especially high priority on the public meaning of a law. Justice Scalia is a frequent advocate of this approach to interpretation, and often resorts to arguments about statutes that are based on what an ordinary person might think a statute means.[20] The general theory behind the argument is that people are entitled to notice of what the law is, so a statute should be taken to mean just what it would mean to an ordinary reader. Letting it mean anything else sets a trap and offends the rule of law. This reasoning is especially powerful in criminal cases, where the defendant's

---

[18] See Manning, *supra* note XX; Adrian Vermeule, Legislative History and the Limits of Judicial Competence: The Untold Story of Holy Trinity Church, 50 Stan. L. Rev. 1833, 1834-35 (1998).

[19] See Frank H. Easterbrook, *Judges as Honest Agents*, 33 Harv. J. L. & Pub. Pol'y. 915 (2010); Abbe R. Gluck, *The States as Laboratories of Statutory Interpretation: Methodological Consensus and the New Modified Textualism*, 119 Yale L.J. 1750 (2010); Adrian Vermeule, *Foreward: System Effects and the Constitution*, 123 Harv. L. Rev. 4, 64-65 (2009).

[20] See Scalia, supra note XX at __; John F. Manning, *Textualism and Legislative Intent*, 91 Va. L. Rev. 419 (2005).

interest in notice—that is, in knowing before one acts what is criminal and what is not—seems especially important.[21] The gun use case is a good example.

This study allows us to suggest another point in favor of asking what ordinary readers would think a statute means, and giving weight to the answer. That question is better than other common questions about meaning at producing answers that aren't contaminated by underlying policy preferences. It may or may not be the question one would most like to have answered about a statute; but a modest question that can be answered relatively well might be better than a perfect question that will tend to be answered badly.

Our second question was whether the respondents to our surveys, even without bias from their policy preferences, were *correct* in their statements about what an ordinary reader would think a statute means? This is surprisingly difficult to answer, even if we assume that those who took our surveys are themselves ordinary readers. It might seem then that we could then look at their views of what these statutes meant, use the results of that inquiry to decide what ordinary readers in fact think, and then compare those findings to what our respondents *predicted* ordinary readers would think. But not so fast. Which of their answers should be used to show what ordinary readers "really" think? We wouldn't want to use everyone's answers to the "ordinary readers" question, because that doesn't show what they thought the statute meant. It shows what they expected *others* to think it meant. We could just look at what our respondents said when they were asked which reading of the statute *they* thought was better. But then we get answers heavily biased by policy preferences—the red line in our graph [assuming we choose that color]. That spoils the inquiry, for a good prediction of what ordinary readers would think of a text isn't supposed to be a prediction of where their biases would lead them. We would end up with a paradox in which opinions about what ordinary readers are valuable because they are unbiased—but also wrong because they are unbiased.

The root problem is that when we ask what an ordinary reader would think a text means, we would like to check the answer against the views of ordinary readers who don't have policy preferences that get in the way of their judgments. It is doubtful whether any such readers are out there. That is one of the implications of this study. One could try setting a baseline by asking some random population of reader what they think a chunk of language means—"using a firearm," perhaps—without any indication of *why* the question is being asked (in other words, without mentioning any legal case). But then the respondents are being forced to interpret the words without a context, and that is a different activity than interpreting them in the particular settings that appeared in our questions. In the end, we suggest that what ordinary readers would think only sounds like an empirical question. It really is not. The ordinary reader is an idealized creature, perhaps not unlike the reasonable person who juries are instructed to imagine in tort cases. Thinking about the

---

[21] See Farnsworth et al., *supra* note XX, at 23; Note, *Textualism as Fair Notice*, 123 Harv. L. Rev. 542 (2009).

ordinary reader is best understood just as a thought experiment, or heuristic. It is a useful device for getting oneself to think a certain way about a text—to focus on the conventional meaning of the words.

IV. LIMITATIONS OF THE STUDY

A. *Causation*

At times in this paper we have spoken of mere "entanglement" between policy preferences and judgments about what a text means. That way of speaking implies no causation. At other points, though, we have talked of policy preferences "influencing" interpretive judgments or having an effect on them. Those claims do suggest causation, of course, so we should consider whether they are hasty. Instead of policy wishes influencing judgments about the text, could judgments about the text somehow be influencing policy wishes?

This is not likely. The policy preferences that respondents display remain the same regardless of what interpretive questions they are asked; but as we have seen, the answers to the interpretive questions sometimes closely follow those preferences and sometimes do not. The causal link we suggest is supported by similar findings in studies of wishful thinking[22] or "halo effects"[23] in non-legal settings. These studies show how underlying preferences about outcomes or similar sources of biases frequently influence judgments about facts, and not the other way around. Our results can be viewed as a particular application of that same general observation.

B. *Criminal cases*

Most of the case studies in our surveys involved federal criminal law. It is possible that the effects found here are special to criminal matters, and would not carry over to civil cases—but again, it is not likely. To address this possibility, we included a non-criminal case that involved an award of attorney's fees at the end of a civil action. As Figure 4 illustrated, respondents displayed the same general pattern in their choices about those cases that they did in the criminal situations, though the effects were somewhat weaker and the pro-government effect surrounding "drafters' intent" does not arise. That last point lends a bit of support to our earlier conjecture that respondents tend to think of the "drafters" as pro-government or more eager to put people in jail. In the civil context, the government is not trying to put someone

---

[22] See Jon Elster, Alchemies of the Mind: Rationality and the Emotions 20-21 (1999); Thomas Gilovich, How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life 75-87 (1991).

[23] See Richard E. Nisbett & Timothy De Camp Wilson, *Telling More Than We Can Know: Verbal Reports on Mental Processes*, 84 Psychol. Rev. 231, 231-32 (1977); Richard E. Nisbett & Timothy De Camp Wilson, *The Halo Effect: Evidence for Unconscious Alteration of Judgments*, 35 J. Personality & Soc. Psychol. 250, 250-52 (1977).

in jail, and it would be surprising to find that respondents thought of the "drafters" as preferring the plaintiff over the defendant, or the other way around.

C. *External validity*

The respondents to most of the questions we have discussed here were students in their first semester of law school. It is possible, of course, that the effects we found are less pronounced among lawyers and judges. But we doubt this on two grounds. First, in prior rounds of this research we did administer our surveys to students at the start and end of their first year of law school. Those surveys did not include the "ordinary readers" question, but they did include other questions considered here—questions about policy preference, about which readings were more consistent with the statute's purpose, and which reading of the statute was better as a matter of text. We found the same results in both populations; there was no significant difference.[24] If a year of law school has not made a dent in the tendency of preferences to influence interpretive judgments, then that tendency is likely stubborn enough to keep exerting some influence later in life. Even if the overall effects were *reduced* in strength, there is no reason to suppose that the relative effects of the different questions we asked would be changed. Asking what an ordinary reader would think of a text would still be a better question than others, even if the benefit in the reduction of bias is less among judges than it is among others.

But in any event we are skeptical about any reduction of these tendencies at all in judges. Consider the relationship between the findings shown here and the following charts, adapted from an earlier empirical study of judicial behavior conducted by one of the authors of this article.[25] Figure 7 is based on career data for all of Supreme Court Justices from 1953-2004.[26] One line shows how often each Justice voted for the government in nonunanimous criminal cases involving constitutional claims. The other line shows their votes in nonunanimous criminal cases depending on some nonconstitutional source of law—usually a statute or rule. The Justices are ordered here according to the data (i.e., by the mean of the two lines) to show the alignment between the two trends.

We can also view the comparison by removing the Justices' names from the graph and instead putting their votes in constitutional cases along the bottom and their votes in nonconstitutional cases along the side. This gives us Figure 8, a scatterplot of the same data that correlates the proportion of the Justices' votes for the government in nonunanimous criminal cases of the two types—constitutional and not. An increase in the share of votes for the government along one of the

---

[24] See Farnsworth et al., *supra* note XX, at ___.

[25] Ward Farnsworth, *Signatures of Ideology: The Case of the Supreme Court's Criminal Docket*, 104 Mich. L. Rev. 67 (2005).

[26] The data were derived from the United States Supreme Court Judicial Database at Michigan State University. For more details and discussion, see Farnsworth, supra note XX, n. 7.

dimensions is very likely to mean an increase along the other. A fitted line shows a strong linear relationship between decisions in favor of the government in either situation.[27]

These charts show that any given Justice votes for the government about as often in cases involving the Constitution as in cases that involve other sources of law. Why should that be? No known theory of interpretation would cause a judge to cast similar votes in cases that depend on entirely different sorts of legal texts. And while originalism, as a constitutional theory, might be expected to produce rulings friendly to the government (because defendants often want the protections of the bill of rights expanded beyond their original meaning), it is hard to see why textualism, intentionalism, or any other approach to interpreting *statutes* would have similar effects. These questions are explored more fully in the earlier work that produced the charts, but the study presented in this article makes a helpful new contribution to an understanding of them.

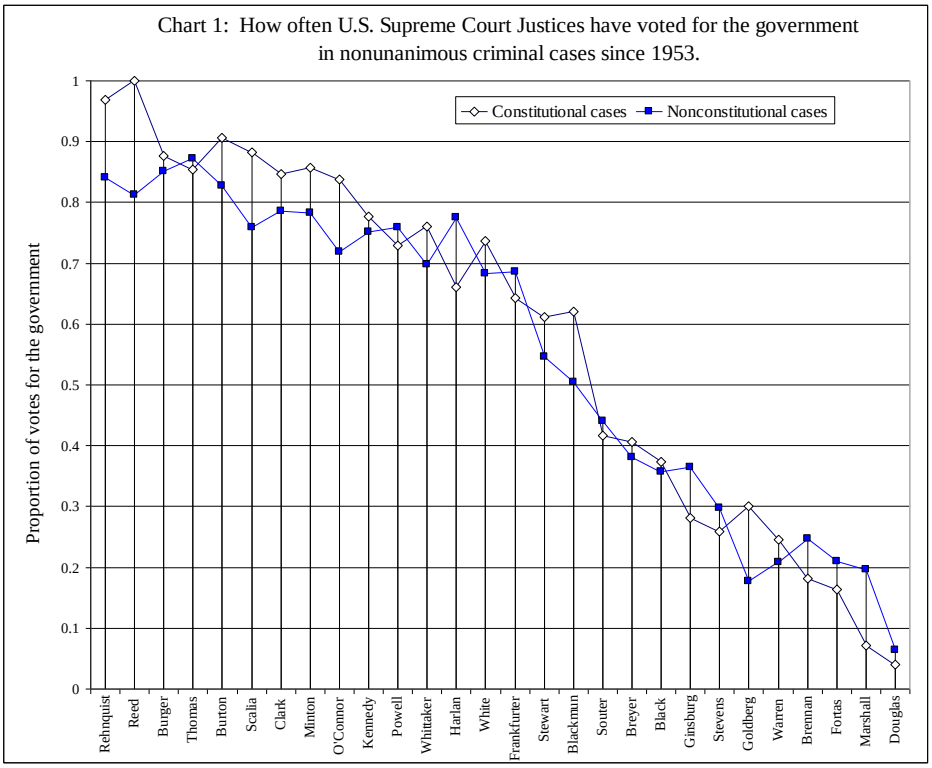

Figure 7

[27] The Pearson Correlation Coefficient (R) is an extremely high .94, accounting for 88% of the variance (R^2).

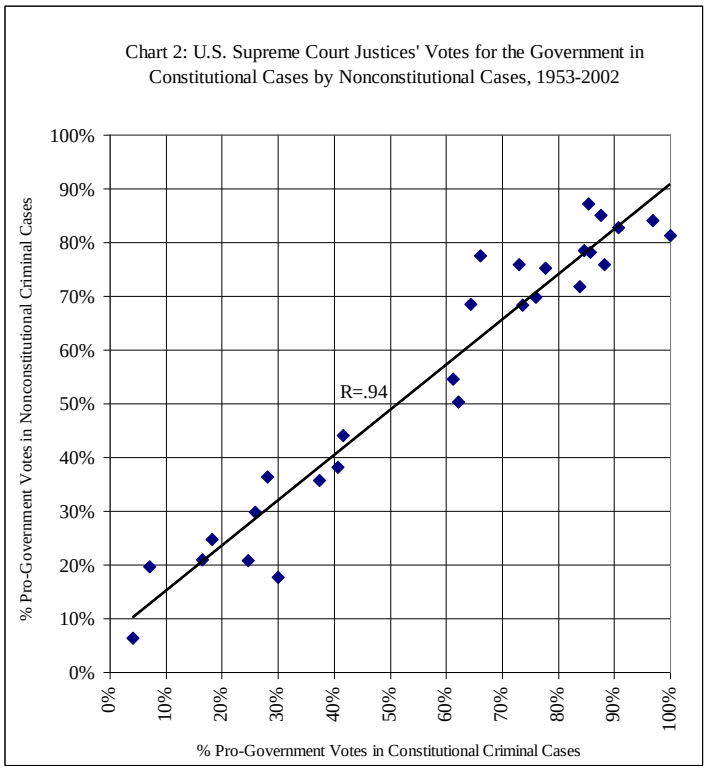

Figure 8

Nonunanimous criminal cases in the Supreme Court are precisely the ones where the legal materials are not conclusive on their face. They contain ambiguities and call for interpretation; most of the cases that served as the basis for questions in our surveys are represented in the set of non-constitutional cases graphed above. Our hypothesis is that when confronted with ambiguous texts—statutes, of course, but probably also cases and constitutional provisions—judges, like other people, have trouble stopping their policy preferences from influencing their judgments. Those policy preferences cut across all sorts of criminal cases, and aren't sensitive to the particular type of legal material (statute, case law, etc.) on which the case seems to depend. In short, when judges vote for the government about as often in close criminal cases of every kind, it is partly because they acting much like the respondents to our surveys.

Obviously this is not a complete explanation of the data just shown. Some judicial votes are better explained in other ways that do not involve the bias exerted by policy preferences. But the purpose of this discussion is not to settle the reasons for judicial dispute. It is to add to the ways that disputes can be explained. The charts just shown are offered here merely to cast doubt on the suggestion that judges are immune to the influences of policy preference that this study has illustrated. The evidence of judicial behavior does not suggest any such immunity. That behavior

shows patterns, rather, that are hard to explain without resort to theories and evidence of the kind we offer here.

CONCLUSION

Policy preferences strongly infect people's judgments about the ordinary meaning of an ambiguous legal text. They also infect judgments about the drafters' intent and predictions of what a court would think the text means. They do not infect nearly as much—sometimes they do not affect at all—judgments about what an ordinary reader of the text would think it means. There are various arguments to be made for and against putting legal weight on that last "objective" question. This Article adds to the arguments in favor of it. The objective question is much easier to answer without bias from policy preferences.

How fully these findings can be applied to judges is an open question, but they help explain a lot of judicial behavior that otherwise is hard to understand. Judges, or for that matter anyone else, can easily decide that they won't let their policy preferences affect their decisions about what a text means. This study suggests that effectively carrying out that decision is harder than it seems. The subjective inner sense that one's preferences are out of the way can be very convincing. That subjective inner sense is a fairly accomplished con artist.

Readers with comments should address them to:

Professor Anup Malani
University of Chicago Law School
1111 East 60th Street
Chicago, IL 60637
amalani@uchicago.edu

Chicago Working Papers in Law and Economics
(Second Series)

For a listing of papers 1–550 please go to Working Papers at http://www.law.uchicago.edu/Lawecon/index.html

551. Douglas G. Baird, Car Trouble, May 2011
552. Omri Ben-Shahar, Fixing Unfair Contracts, May 2011
553. Saul Levmore and Ariel Porat, Bargaining with Double Jeopardy, May 2011
554. Adam B. Cox and Richard T. Holden, Reconsidering Racial and Partisan Gerrymandering, May 2011
555. David S. Evans, The Antitrust Economics of Free, May 2011
556. Lee Anne Fennell, Property and Precaution, June 2011
557. Omri Ben-Shahar and Anu Bradford, Reversible Rewards, June 2011
558. Alon Harel and Ariel Porat, Commensurability and Agency: Two Yet-to-Be-Met Challenges for Law and Economics, June 2011
559. Randal C. Picker, After Google Book Search: Rebooting the Digital Library, June 2011
560. Julie A. Roin, Privatization and the Sale of Tax Revenues, July 2011
561. Joseph Issenbergh, Last Chance, America, July 2011
562. Richard H. McAdams, Present Bias and Criminal Law, July 2011
563. David Weisbach, Is Knowledge of the Tax Law Socially Desirable?, July 2011
564. Louis Kaplow and David Weisbach, Discount Rates, Judgments, Individuals' Risk Preferences, and Uncertainty, July 2011
565. Louis Kaplow, Elisabeth Moyer and David A. Weisbach, The Social Evaluation of Intergenerational Policies and Its Application to Integrated Assessment Models of Climate Change, July 2011
566. David A. Weisbach, Carbon Taxation in Europe: Expanding the EU Carbon Price, July 2011
567. Nuno Garoupa and Tom Ginsburg, Hybrid Judicial Career Structures: Reputation v. Legal Tradition, August 2011
568. Richard A. Epstein, F. Scott Kieff, and Daniel F. Spulber, The FTC's Proposal for Regulationg IP through SSOs Would Replace Private Coordination with Government Holdups, August 2011
569. Dhammika Dharmapala, Nuno Garoupa, and Richard McAdams, Do Exclusionary Rules Convict the Innocent? August, 2011
570. Andres Sawicki, Better Mistakes in Patent Law, August 2011
571. Jonathan S. Masur and Eric A. Posner, Regulation, Unemployment, and Cost-Benefit Analysis, August 2011
572. Adam B. Cox and Eric A. Posner, Delegation in Immigration Law, September 2011
573. Joseph Bankman and David Weisbach, A Critical Look at a Critical Look—Reply to Sanchirico, September 2011
574. M. Todd Henderson and Frederick Tung, Pay for Regulator Performance, September 2011
575. William H. J. Hubbard, The Problem of Measuring Legal Change, with Application to *Bell Atlantic v. Twombly*, September 2011
576. Adam M. Samaha, Regulation for the Sake of Appearance, October 2011
577. Ward Farnsworth, Dustin Guzior, and Anup Malani, Implicit Bias in Legal Interpretation, October 2011