Maurer School of Law: Indiana University

Digital Repository @ Maurer Law

Articles by Maurer Faculty

Faculty Scholarship

2017

After the Override: An Empirical Analysis of Shadow Precedent

Deborah A. Widiss

Indiana University Maurer School of Law, dwidiss@indiana.edu

Brian J. Broughman

Indiana University Maurer School of Law, bbroughm@indiana.edu

Follow this and additional works at: <https://www.repository.law.indiana.edu/facpub>

Part of the [Judges Commons](#), and the [Legislation Commons](#)

Recommended Citation

Widiss, Deborah A. and Broughman, Brian J., "After the Override: An Empirical Analysis of Shadow Precedent" (2017). *Articles by Maurer Faculty*. 2545.

<https://www.repository.law.indiana.edu/facpub/2545>

This Article is brought to you for free and open access by the Faculty Scholarship at Digital Repository @ Maurer Law. It has been accepted for inclusion in Articles by Maurer Faculty by an authorized administrator of Digital Repository @ Maurer Law. For more information, please contact rvaughan@indiana.edu.

LAW LIBRARY
INDIANA UNIVERSITY
Maurer School of Law
Bloomington

After the Override: An Empirical Analysis of Shadow Precedent

Brian J. Broughman and Deborah A. Widiss

ABSTRACT

Congressional overrides of prior judicial interpretations of statutory language are typically defined as equivalent to judicial overrulings, and they are presumed to play a central role in maintaining legislative supremacy. Our study is the first to empirically test these assumptions. Using a differences-in-differences research design, we find that citation levels decrease far less after legislative overrides than after judicial overrulings. This pattern holds true even when controlling for depth of the superseding event or considering only the specific proposition that was superseded. Moreover, contrary to what one might expect, citation levels decrease more quickly after restorative overrides—in which Congress repudiates the prior Supreme Court decision as incorrect—than after overrides intended to update or clarify the law. This suggests that ongoing citation of overridden precedents, what we call shadow precedents, may be driven more by information failure or ambiguity than by ideological disagreements between the branches of government.

1. INTRODUCTION

The ability of Congress to override judicial interpretations of statutes is central to theories of the separation of powers. While the Constitution formally places all law-making authority in Congress, judicial decisions informally shape legislation by filling in gaps and resolving ambiguity.

BRIAN J. BROUGHMAN is Associate Dean for Research, Professor of Law, and Director of the Center for Law, Society, and Culture at Indiana University. DEBORAH A. WIDISS is a Professor of Law at Indiana University. We are very grateful to Matthew Christiansen, Bill Eskridge, and James Spriggs for their generosity in sharing data with us. For helpful comments on this project, we thank Lynn Baker, Bill Eskridge, Michael Frakes, Michael Gilbert, Abbe Gluck, Mitu Gulati, David Hamilton, Rick Hasen, William Hubbard, Dan Klerman, Maggie Lemos, Tim Meyer, Barry Pyle, Neil Siegel, James Spriggs, Matthew Stephenson, Stephen Wasby, Abby Wood, an anonymous referee, and workshop participants at Duke Law School, Indiana University's Maurer School of Law, the University of Southern California's Gould School of Law, the University of Texas School of Law,

[*Journal of Legal Studies*, vol. 46 (January 2017)]

© 2017 by The University of Chicago. All rights reserved. 0047-2530/2017/4601-0002\$10.00

Legislative supremacy thus depends on the assumption that, if Congress disagrees with a judicial interpretation of a law, it may override that interpretation by passing a new statute or amending an existing statute (Barnes 2004; Eskridge 1994; Levi 1949).

Accordingly, legislative overrides play a large role in both political science and legal scholarship. Positive political theorists contend that overrides constrain the Supreme Court's ability to follow its own ideological preferences (Ferejohn and Weingast 1992; Gely and Spiller 1990; Bergara, Richman, and Spiller 2003; Bailey and Maltzman 2011), whereas legal theorists more typically present overrides as helpful corrections from Congress (for example, Eskridge 1994; Elhauge 2002; Marshall 1989). But both accounts depend on two conditions being satisfied (Widiss 2009). First, Congress must monitor statutory interpretation decisions and respond to decisions with which it disagrees. Empirical studies show that Congress, while limited by gridlock in recent years (Hasen 2013), does regularly enact overrides (Eskridge 1991a; Klerman 2007; Staudt, Lindstädt, and O'Connor 2007; Hasen 2013; Christiansen and Eskridge 2014). Second, congressional overrides must have some bite—they must supersede the prior judicial interpretation. This is generally assumed but not addressed. Ours is the first empirical study to measure the extent to which an override changes citation patterns to the overridden case.¹

A congressional override is typically defined as the legislative equivalent of judicial overruling. However, on the ground for the lower courts that must first interpret the significance of a change in the law, they are quite different. If a decision has been overruled by a higher court, it is clear to lower courts that they should follow the signals of that higher court. In addition, in most instances, such changes are immediately flagged by legal research tools like Westlaw and LexisNexis. By contrast, it often takes several years for Westlaw and LexisNexis to indicate that a new statutory provision affects the validity of a precedent (Widiss 2014; Christiansen and Eskridge 2014). Even if aware of the override, lower

Washington University School of Law, the 2015 American Law and Economics Association annual meeting, the 2015 Midwest Political Science Association annual meeting, and the 2016 Legislation Scholars Roundtable held at Cardozo Law School. We would also like to thank Sarah Armstrong, Stacey Kaiser, Lisa Moat, and Lyndsey Mulherin for valuable research assistance; we owe a particular debt of gratitude to Matt Pfaff, who went far beyond the typical research assistant role in his contributions to design and implementation of this study.

1. A few studies have assessed other aspects of courts' implementation of override statutes (Barnes 2004; Christiansen and Eskridge 2014).

courts must determine the extent to which the new statutory language supersedes the rule established by the precedent. Faced with competing signals from Congress and the Supreme Court, a lower-court judge may be apt to continue to follow the superseded precedent, at least on any point where it is at all unclear whether it remains controlling. This could be the result of a generalized deference to higher courts or a preference for reducing the risk of reversal by deciding cases in line with the Supreme Court's presumed preferences. The ambiguity implicit in interpreting overrides also might offer lower-court judges the opportunity to advance their own ideological preferences. As a result, even after Congress passes an override, the precedent may live on as what we call a shadow precedent. Earlier qualitative work identifies numerous examples of courts following shadow precedents in employment discrimination decisions (Widiss 2009, 2012, 2015). This project examines the extent to which this is a more general phenomenon.

The theory of shadow precedent predicts that, everything else equal, an overridden case is more likely than an overruled case to be treated as valid precedent after the superseding event. To investigate this question, we put together an original database of judicial citations to three different groups of Supreme Court decisions: cases overridden by Congress ($n = 166$), cases subsequently overruled by the Court itself ($n = 55$), and a matched control group, created using coarsened exact matching (CEM) (Blackwell et al. 2009), of Supreme Court decisions that were neither overridden nor overruled ($n = 141$). For each case in our data set, we record the number of annual citations to the case, sorted by Shepard's signal, for a 16-year period, starting 5 years prior to the event (override or overruling) and continuing until 10 years after the event. We use a differences-in-differences research design with a matched control group to compare how the case was cited before and after the override or overruling.

Both the overruled cases and the overridden cases receive more negative citations postevent than the control group. However, there are important differences. For judicially overruled cases, the negative citations quickly become more common than positive citations, and the total number of citations falls dramatically. By contrast, for legislatively overridden cases, the number of positive citations and the overall number of citations show little change. Even 10 years after an override has been enacted, most overridden precedents are still widely cited as controlling law.

To provide a more nuanced comparison, we assess the depth of the

superseding event (that is, how fully the override or overruling departed from the precedent). We find that for both groups of cases, depth of the superseding event is associated with fewer postevent citations, but at every level of depth, citations to cases overruled by a judicial opinion decline more than citations to cases overridden by Congress. We find similar results when we consider the explicitness with which the Court or Congress expressed intent to supersede a prior decision. For a randomly selected subset of overridden and overruled cases, we conduct more fine-grained analysis by using Lexis-Nexis headnotes to isolate the legal proposition directly affected by the superseding event and compare them to unrelated headnotes for the same case. We find that for both sets of cases, directly superseded headnotes receive significantly fewer citations than unrelated headnotes, but the decline is more substantial for overruled cases. And finally, we find that the number of citations drops more quickly after restorative overrides—which repudiate the prior judicial decision as contrary to congressional intent (Christiansen and Eskridge 2014)—than after overrides that are intended to update or clarify a law. This suggests that ongoing citation of overridden cases may be driven primarily by information failure or ambiguity rather than ideological fights between the branches of government.

In sum, on average, looking at specific superseded headnotes and using regression analysis to control for relevant factors such as depth and explicitness, we find that the precedential value of a superseded case dissipates far more quickly and completely after a judicial overruling than after a congressional override. Our findings are robust to alternative econometric specifications. We control for numerous considerations that may affect postevent citations, including ideology, characteristics of cases and overrides, and the inclusion of year and case fixed effects. We include subsample analysis showing that our results also apply to alternative measures of precedent, different event windows, and a balanced panel. While legislative overrides and judicial overrulings are not exogenous events, our use of case fixed effects, inclusion of a matched control group, and separate analysis of headnotes reduces concerns associated with unobserved effects.

Our findings are consistent with the theory of shadow precedent: legislatively overridden cases are more likely to continue to be cited than judicially overruled cases. This is contrary to the conventional view that overrides are functionally equivalent to overrulings. Our results further suggest that information failure and ambiguity are likely causes of this

ongoing reliance on overridden precedents and that these factors are more important than ideological differences between the Court and Congress. However, additional empirical research and theoretical refinement may be necessary to fully explain the differences we observe in citation patterns between overridden and overruled cases. That said, our core empirical result is noteworthy in itself, regardless of the reason for the divergent citation patterns, as it suggests a need to rethink theories of statutory interpretation that rely on congressional overrides to redirect judicial interpretations. In short, our findings suggest that overrides are not fully serving the role they are expected to play in ensuring legislative supremacy. This is particularly important for overrides that are intended to update or clarify statutory law—and that are often enacted by Congress in response to judicial invitation—but for which we find almost no effect on the level of citations to the preexisting case.

The rest of our paper is organized as follows: Section 2 surveys the background literature on precedent and legislative overrides and develops testable predictions. Section 3 describes our data, Section 4 tests the shadow-precedent theory using fixed-effects regression analysis and includes a number of robustness checks, and Section 5 discusses the implications of our research and concludes.

2. BACKGROUND LITERATURE AND THEORY

This section begins with an overview of existing research discussing the extent to which precedent and overrides are potential constraints on judicial behavior. After an override, these constraints are in tension with each other: the precedent will pull in one direction, and the text of the override will pull in another. Lower courts are caught in the middle, as they are asked to resolve this tension with little guidance from Congress or the Court. We end this section with testable predictions regarding the effect of overrides on precedents.

2.1. Background Literature

Adherence to precedent is a central foundation of the American judicial system. In general, courts are expected to decide relevantly similar cases consistently, which promotes efficiency, fairness, and predictability (Lindquist and Cross 2005; Schauer 1987). Some empirical studies of Supreme Court decisions find that ideological preferences play a large role

in decisions and that precedent, by contrast, offers comparatively little constraint (for example, Segal and Spaeth 2002), whereas other studies find that precedent matters in a variety of contexts (for example, Bailey and Maltzman 2011; Richards and Kritzer 2002; Hansford and Spriggs 2006). These findings are likely shaped in part by docket selection; the Supreme Court generally takes cases for which there has been a circuit split and thus for which, almost by definition, existing precedent does not clearly establish the proper outcome (Cross 1997).

Research on decisions by lower-court judges tells a somewhat different story. Both district court and circuit court judges generally comply with Supreme Court precedent that is clearly on point (for example, Kim 2007; Klein 2002; Songer and Sheehan 1990). Where application of a precedent is unclear, however, studies suggest that judges' own ideology (Boyd and Spriggs 2009; Sunstein et al. 2006), their network of peer judges (Choi and Gulati 2008; Choi, Gulati, and Posner 2012), the composition of the panel with whom they sit (Sunstein et al. 2006; Kim 2009), the presumed ideological preferences of reviewing courts (Randazzo 2008; Westerland et al. 2010), and changes in personnel on the Supreme Court (Benjamin and Vanberg 2016) all may play a role. There has been relatively little research into how lower courts implement Supreme Court decisions that overrule earlier decisions. Benesh and Reddick (2002) find that lower courts respond quickly to such changes in the law,² whereas Tokson (2015), which considers changes initiated by the Supreme Court and statutory changes, finds that lower courts sometimes fail to adopt fully the new doctrine, particularly if the new regime is difficult to apply or replaces a rule with a standard. Our study adds not only to the understanding of the effects of an override but also to this literature on the effects of a judicial overruling.

In the realm of statutory interpretation, the possibility of congressional override is typically presented as a significant additional limitation on courts. Positive political science models often posit that the Supreme Court will interpret a statute in a manner that is as close to its ideological preferences as possible without triggering a legislative override of the decision (for example, Ferejohn and Weingast 1992; Gely and Spiller 1990; Eskridge 1991b). Empirical studies are mixed, with some finding evidence that the Court, at least in some instances, is constrained by the possibil-

2. This accords with studies that find lower courts to be responsive to significant changes in the doctrine governing particular areas of law (for example, Luse et al. 2009; Songer, Segal, and Cameron 1994).

ity of an override (Spiller and Gely 1992; Bergara, Richman, and Spiller 2003; Bailey and Maltzman 2011; Pacelle, Curry, and Marshall 2011) and others finding that the Court generally rules according to its ideological preferences, without adjusting its behavior to avoid a response from Congress (Segal 1997). Traditional legal theory, by contrast, typically conceives of overrides as part of a conversation between the courts and Congress, in which courts interpret statutes in line with established legal principles and welcome corrections by Congress if they misunderstand congressional intent or if the policy needs to be updated (for example, Marshall 1989; Levi 1949). The Supreme Court also frequently announces this understanding of the role of overrides (for example, *Flood v. Kuhn*, 407 U.S. 258 [1972]), and it regularly invites Congress to override decisions (Christiansen and Eskridge 2014).

These legal and positive political theories, as well as the rationales espoused in Supreme Court doctrine, depend on two assumptions: (1) that Congress monitors judicial opinions and enacts overrides when necessary to correct or update statutory policy and (2) that enactment of an override will change subsequent judicial behavior.³

Several studies examine the validity of this first proposition by seeking to catalog all statutory provisions that supersede prior statutory interpretation decisions by the courts. This work establishes that overrides are fairly common; they occur in virtually all areas of federal statutory law, but they are especially prevalent in federal procedure, civil rights, tax, criminal, and bankruptcy (Eskridge 1991a; Hausegger and Baum 1998; Staudt, Lindstädt, and O'Connor 2007; Hasen 2013; Christiansen and Eskridge 2014; Buatti and Hasen 2015; Christiansen, Eskridge, and Thypin-Bermeo 2015).⁴

There has been very little consideration, however, of the second question: what happens after an override? In other words, are overrides effective in changing the law as applied on the ground? This poses two distinct

3. This second assumption is often only implicit, but positive political science models positing that overrides serve as a constraint on judicial interpretation obviously assume that enactment of an override would curtail subsequent judicial interpretation. Similarly, the conversation between the judiciary and Congress that legal scholars imagine would be ineffective if Congress's half of the conversation goes unheeded.

4. Within this literature, there is disagreement as to whether overrides should be defined to include all statutory provisions that modify the result in prior statutory interpretation decisions (Christiansen and Eskridge 2014; Christiansen, Eskridge, and Thypin-Bermeo 2015) or whether the category should be limited to conscious overrides (Hasen 2013; Buatti and Hasen 2015). We discuss these definitional issues and how they affect our study below and in the online appendix.

questions. The first is how override statutes are interpreted; the second is how the enactment of an override changes reliance on overridden cases. To our knowledge, there have been two quantitative studies—Barnes (2004) and Christiansen and Eskridge (2014)—that explore aspects of how courts interpret override statutes but none that addresses the effect of an override on precedent.

Ours is the first large-scale quantitative study of how enactment of an override changes reliance on the overridden case—or what we term a shadow precedent. While it may seem counterintuitive that courts would cite overridden precedents at all, earlier qualitative work provides examples of this phenomenon in the employment discrimination context. Widiss (2009) shows that courts sometimes continue to follow the rationale or reasoning supporting a holding that has been superseded, on the grounds that the override statute addresses only an application of that reasoning.⁵ Widiss (2012) documents lower courts' confusion when Congress enacts an override but does not amend the other statutes with similar language that have typically been interpreted consistently.⁶ In the ex-

5. For example, in *General Electric Co. v. Gilbert* (429 U.S. 125 [1976]), the Supreme Court held that the prohibition on discrimination on the basis of sex found in Title VII of the Civil Rights Act of 1964 did not prohibit discrimination on the basis of pregnancy, reasoning that the policy at issue distinguished between “pregnant women and nonpregnant persons” rather than between men and women. Two dissenting opinions argued that pregnancy discrimination was inherently a form of sex discrimination, since only women become pregnant. Congress quickly superseded *Gilbert* by enacting the Pregnancy Discrimination Act of 1978, which amended Title VII to define sex as including but not limited to “pregnancy, childbirth, or related medical conditions” (Pub. L. No. 95-555, 92 Stat. 2076). In more recent years, women have alleged that denial of access to contraceptives or discrimination because of lactation are forms of discrimination because of sex. Because these contexts (arguably) are not directly addressed by the language of the Pregnancy Discrimination Act, some lower courts have followed *Gilbert* as controlling precedent (see, for example, *Martinez v. NBC*, 49 F. Supp. 2d 305, 309 [SDNY 1999]), while others have followed the reasoning from the dissent in *Gilbert* (see, for example, *Erickson v. Bartrell Drug Co.*, 141 F. Supp. 2d 1266, 1270 [W.D. Wash. 2001]). In recent years, several courts have held that lactation is encompassed within the “related medical condition” provision of the statutory language (see, for example, *EEOC v. Houston Funding II, Ltd.*, 717 F.3d 425 [5th Cir. 2013]).

6. This has been widely litigated in the context of the standard for causation under various employment discrimination statutes. In 1991, Congress amended Title VII of the Civil Rights Act of 1964 to supersede a prior Supreme Court decision (*Price Waterhouse v. Hopkins*, 490 U.S. 228 [1989]) regarding the causation standard that governs claims of discrimination under Title VII. Congress did not, however, amend other employment discrimination statutes, such as the Age Discrimination in Employment Act, the Americans with Disabilities Act, or a distinct part of Title VII, even though all of these contexts had typically been interpreted consistently. Lower courts divided over what causation standard should apply to these other contexts (Widiss 2009). In *Gross v. FBL Financial*

amples above, courts and commentators could reasonably disagree with the propriety of continuing to follow the overridden precedent. Widiss (2015) shows that courts also sometimes simply make mistakes, applying statutory standards that have unquestionably been superseded.⁷ Whether because of disagreement over the scope of an override, ideological preferences, confusion, ignorance, or simply resistance to change (compare Tokson 2015), overridden precedents remain influential.

2.2. Shadow Precedent: Theory and Predictions

Both judicial overrulings and legislative overrides are intended to supersede, at least in part, the rule established in a prior decision. It is common, in fact, to define overrides as the legislative equivalent of overruling. However, there are differences between overrulings and overrides that suggest that overrides may be less effective than overrulings at changing citation patterns to the precedent.

As a preliminary matter, after either an overruling or an override, courts must recognize that something has occurred that could affect precedential value of the prior case. Courts and lawyers rely heavily on Westlaw, LexisNexis, and other legal research services to flag when subsequent developments affect the reliability of a precedent, either positively or negatively. As described in Widiss (2014), the legal research services have adopted coding protocols that look almost exclusively to judicial signals. Thus, when the Supreme Court explicitly overrules a prior Supreme Court decision, both Westlaw and Lexis immediately “red flag” the prior decision. By contrast, a case generally will not be identified as superseded by a statute until a court issues a decision that makes this connection; consequently, there is often a multiyear lag before legal research

Services (557 U.S. 167 [2009]) and *University of Texas Southwestern Medical Center v. Nassar* (133 S. Ct. 2517 [2013]), the Supreme Court instructed lower courts that the causation standard specified by the override should not be applied in these other contexts.

7. Widiss (2015) illustrates this phenomenon by looking at the implementation of the Americans with Disabilities Act Amendments Act of 2008 (ADAAA) (Pub. L. No. 110-325), an unusually strong and clear override. The ADAAA includes statutory purposes clauses that explicitly reject the Supreme Court’s prior interpretations of the Americans with Disabilities Act as counter to Congress’s original intent; the ADAAA also amended the substantive language of the Americans with Disabilities Act to supersede the Court’s prior interpretations. Nonetheless, as documented in Widiss (2015), lower courts regularly continue to cite the overridden precedents as controlling law. Some of these lower-court decisions make no mention of the ADAAA; others acknowledge that the ADAAA was enacted but fail to recognize that the new statutory language unquestionably supersedes relevant portions of the precedent.

services recognize an override (Widiss 2014). These lag times suggest that information failure may contribute to ongoing reliance on overridden precedents.

Once aware of the superseding event, lower courts must determine how it affects the precedential value of the prior decision. Again, with judicial overrulings, this is usually relatively straightforward in that a court needs only to parse the overruling decision. After an override, the analysis is often more complex: lower courts must determine the extent to which the new statutory language supersedes the precedent, which otherwise remains binding on lower courts. In resolving this tension, lower-court judges might lean toward following the precedent, at least when it is ambiguous which should control. This could be for abstract rule-of-law reasons or for more instrumental reasons. That is, for a trial court judge, the possibility of review and potential reversal by an appellate court or the Supreme Court (the source of the precedent and the judge's superiors in the judicial hierarchy) is likely to be of more immediate concern than any hypothetical feedback from a future Congress (the source of the override). It is also possible that courts use the ambiguity implicit in overrides to advance their own ideological preferences. For these reasons, as well as potential information failure as discussed above, the theory of shadow precedent predicts that, everything else being equal, an overridden case is more likely than an overruled case to be cited as valid precedent after the respective event (the shadow-precedent hypothesis).

We use distinctions between restorative and nonrestorative overrides, as classified by Christiansen and Eskridge (2014), to assess further the factors that may drive ongoing citation of overridden precedents. Restorative overrides, defined as overrides that repudiate a prior judicial decision as contrary to congressional intent, tend to be very explicit, and the fight between Congress and the Court often receives significant coverage in the legal and popular presses. By contrast, nonrestorative overrides that update or clarify the law, like major revisions of bankruptcy law or the tax code, can be quite deep—in that they wholly replace one or several precedents—but they are less likely to denigrate, or even identify, the precedents affected. If ongoing reliance on shadow precedents stems primarily from information failure or from the failure of Congress to give clear signals, precedent superseded by a restorative override will be less likely to be cited positively after an override than precedent superseded by a nonrestorative override.

On the other hand, restorative overrides occur more frequently in

areas of the law where there are sharp partisan divides. In addition, the fact that Congress is so clearly disagreeing with the Court could increase the likelihood that lower courts would feel pressure to interpret an override as narrowly as possible, in that they can reasonably predict that a majority of the Supreme Court would prefer a different interpretation than that which Congress has enacted. To the extent that lower courts' compliance with the assumed preferences of the Supreme Court drives ongoing citations to shadow precedents, precedent superseded by a non-restorative override would be less likely to be cited positively than precedent superseded by a restorative override.

3. DATA

To investigate these questions, we constructed a database of citations to Supreme Court decisions. Citation counts are a common mechanism to gauge the precedential importance of a case (see, for example, Black and Spriggs 2013; Hansford and Spriggs 2006; Westerland et al. 2010). Our database includes annual citations to 166 statutory interpretation cases subsequently overridden by Congress, 55 cases subsequently overruled by the Court, and a matched control group of 141 Supreme Court decisions that were neither overridden nor overruled.⁸

For the sample of cases that were overridden (hereafter, the overridden sample), we collect data for all cases (decided after 1946)⁹ identified by Christiansen and Eskridge (2014) as being subject to legislative overrides enacted between 1985 and 2011 ($n = 166$). The definition of an override in Christiansen and Eskridge (2014)—any statutory provision that modifies the result in a prior statutory interpretation decision—is broader than in Buatti and Hasen (2015), which includes only conscious overrides (where the legislative history or statutory language makes clear that Congress was responding to a judicial decision). We use the former definition because, under standard legal principles, applicable statutory language should govern resolution of cases, whether or not the interaction of that statutory language with a precedent was clearly identified in legislative history. Nonetheless, we acknowledge that the consciousness

8. See our online data files. The online appendix includes STATA code for replicating the regression results included in the tables.

9. This limitation comes from the fact that we gather background data on the cases from the Supreme Court Database, which includes all Supreme Court decisions after 1946 (see Washington University Law, The Supreme Court Database [<http://scdb.wustl.edu>]).

of congressional action may affect how precedent is cited, and thus we include *Conscious* as an explanatory variable in the analysis below.

To compile the sample of cases that were overruled (hereafter, the overruled sample), we use Brenner and Spaeth (1995) and the Supreme Court Database (SCD) to identify cases (decided after 1946) that were overruled between 1985 and 2011 ($n = 55$). It would be ideal to compare only statutory interpretation cases within the two categories. However, because it is relatively uncommon for the Court to overrule prior statutory decisions, our list of overruled cases also includes constitutional decisions.

Finally, we use CEM (Blackwell et al. 2009) to construct a contemporaneous control group of decisions that were neither overridden nor overruled. In CEM, a variant of exact matching, the data are first coarsened into categories defined by the researcher and then exact matched using the coarsened data. This process improves “estimation of causal effects by reducing imbalance in covariates between treated and control groups” (Blackwell et al. 2009, p. 524). Using data available in the SCD, treatment group cases are matched on the basis of six observed characteristics: year of the decision (coarsened by 2 years), ideological direction (liberal or conservative), area of law (divided into 21 categories), type of law (statutory, constitutional, or other, such as court rules or diversity cases), type of decision (signed opinions, judgments [plurality opinions], or per curiam opinions), and number of votes for the majority opinion.¹⁰ We found a 1:1 match for 102 cases from the overridden group and 39 cases from the overruled group, which gives us a total of 141 matched control group cases. Because our matched control group covers the same time period, general subject area, and ideology as the two treatment groups, it can help us isolate the effect of treatment as opposed to unobserved developments occurring within our event window.

For each case in our sample, we collected the number of annual citations and associated Shepard’s signals for the 16-year period starting 5 years prior to the event and continuing until 10 years after the event. Going back 5 years prior to the event gives us a solid baseline of how each case is cited before the legislative override or judicial overruling. Because overrides are not retroactively implemented, we use a longer postevent period—10 years—to capture the full impact of the superseding event. We treat this 16-year period as an event window indexed by t , from $t =$

10. Majority votes are divided into three categories: four or five votes, six or seven votes, and eight or nine votes.

–5 to $t = 10$ and with the event (override or overruling) centered at $t = 0$. This effectively gives us panel data with up to 16 observations per case, with the case-year pair as the relevant unit of analysis.

Table 1 provides descriptive statistics. All of the overrides and overrulings occurred between 1985 and 2011; the average year for both groups was approximately 1995. The average age of the cases that were superseded by these events, by contrast, varies, with the mean case in the overridden group decided in 1986, and the mean case in the overruled case decided in 1973. Relatedly, the average amount of time from decision until override (8.4 years) is much shorter than from decision until overruling (22.6 years). This reflects a difference between overrides and overrulings. The Court will not typically overrule its own precedents unless there has been some significant intervening development that can plausibly justify abandoning *stare decisis* principles. By contrast, Congress often acts very quickly to supersede judicial decisions with which it disagrees; 27 percent of the cases in our overridden group were superseded less than 2 years after the decision.¹¹ Our use of the control group, which matches the age of cases in the overruled and overridden groups, helps us distinguish changes in citation levels that are responsive to the superseding event from the more general depreciation—that is, gradual decline in citations—that affects all precedents (see, for example, Landes and Posner 1976; Merryman 1954; Black and Spriggs 2013).

We use the SCD's classifications of cases as liberal or conservative as a rough gauge of the ideological directions of the decisions. A significantly higher percentage of cases in the overruled group are classified as liberal decisions (65 percent) than in the overridden group (43 percent), which reflects the changing composition of the Court and Congress over this period.

Table 1 also reports the average number of citations that each case received per year. On average, we were able to collect 13.7 years of citation data for cases in the overridden sample, 14.5 years for cases in the overruled sample, and 14.2 years for cases in the control group. For all three groups, the mean number of annual citations to each case is substantially

11. This results in an unbalanced panel, with pre-event observations for some overridden cases truncated by the amount of time between the decision and the superseding event; a shorter window, however, may result in a less accurate baseline measure for pre-event citations. In Section 4.3, we explore the robustness of our analysis using a balanced panel.

Table 1. Summary Statistics

	Overridden (n = 166)			Overruled (n = 55)			Matched Control (n = 141)		
	Mean	Median	SD	Mean	Median	SD	Mean	Median	SD
Overview:									
Event year	1994.8	1995	6.63	1995.5	1994	8.07			
Decision year	1986.4	1988	11.02	1972.9	1974	13.20	1982.9	1985	13.44
Decision to event (years)	8.40	4	10.46	22.55	21	12.50			
<2 (%)	.27		.44	.00		.00			
2–10 (%)	.46		.50	.20		.40			
>10 (%)	.27		.45	.80		.40			
Majority votes	6.66	6	1.53	6.38	6	1.39	6.88	6	1.53
Minority votes	2.14	2.5	1.53	2.36	3	1.47	1.94	2	1.52
Liberal case	.43			.65			.54		
Case name in override	.07								
Restorative	.22								
Updating or clarifying	.78								
Citations:									
Years of citation data per case	13.66	14	2.49	14.50	16	2.42	14.19	15	2.25
Total citations per year	56.97	25.69	142.70	128.57	22.07	385.84	40.06	12.08	92.39
Net citations per year	52.93	24.08	141.37	114.22	16.62	349.02	38.07	10.77	90.24
Shadow-precedent score	.89	.71	.62	.58	.56	.41	.97	.86	.81

greater than the median, which reflects outlier cases in each group.¹² According to medians, a typical case from the overridden and the overruled groups receives a similar number of annual citations (25.7 and 22.1, respectively), while the matched control group case receives fewer (12.1).

Using Shepard's signal indicators, we measure net citations to each case in year t as

$$\text{Net Citations}_t = (\text{Positive} + \text{Neutral} + \text{Cited By}) \\ - (\text{Warning} + \text{Caution} + \text{Questioned}).$$

The most common Shepard's signal is "cited by." Signals indicating more extensive discussion, such as positive treatment or warning, are comparatively rare. We include cited-by signals in our calculation of net citations since even such neutral signals indicate that later courts cited the prior case as presumptively valid precedent. However, as described below, we test the robustness of our results against alternative methods of citation counting that give more weight to small fluctuations in negative citations and a variation of this measure that excludes neutral and cited-by cites entirely.

Our primary interest is not in the absolute (or even net) number of citations that a case receives per year but rather the change in the number of citations that accompanied the event. Did net citations decline following the event, and if so how big was the change? To provide a rough case-level measure of this, we assign each case a shadow-precedent score, defined as

Shadow-Precedent Score

$$= \frac{\text{Average net citations per year in the postevent period}}{\text{Average net citations per year in the pre-event period}},$$

where the pre-event period is from year $t = -5$ to year $t = -1$ and the postevent period is from year $t = 3$ to year $t = 10$. We exclude years immediately following the passage of the override or the overruling ($t = 0, 1$, and 2); this is because overrides are generally prospective only, and thus claims litigated during this period may still be adjudicated under

12. For example, the overruled group includes *Conley v. Gibson* (355 U.S. 41 [1957]), which is cited extremely frequently because it addresses the standard for a motion to dismiss, and the overridden group includes several habeas corpus cases commonly referenced in (the voluminous body of) prisoner litigation. In the regression analysis below, our dependent variable is defined to reduce the impact of heavily cited cases.

the old statutory language.¹³ To give us a more accurate baseline rate of citations, we exclude observations from the year in which the case was decided and cases that received fewer than three citations per year in the pre-event period. As a result of these restrictions, shadow-precedent score is defined for 132 cases in the overridden sample, 49 in the overruled sample, and 113 in the control group.

Consistent with our hypothesis, the mean shadow-precedent score is significantly higher for cases in the overridden sample than for cases in the overruled sample. In the years following an override, we find that an average overridden case typically receives 89 percent (median = 71 percent) of the number of annual net citations compared with the same case in the years prior to the override. By contrast, we find that overruled cases experience a significantly larger drop in citations after the event, falling on average to 58 percent (median = 56 percent) of the pre-event level of net citations. The average shadow-precedent score for cases in the control group is 97 percent (median = 86 percent).

Table 2 shows the average shadow-precedent score sorted by the depth of the superseding event. We use the Christiansen and Eskridge (2014) classification of depth, a scale of 1 to 5, for cases in the overridden group; we create a similar classification for cases in the overruled group.¹⁴ While there is considerable variation, for both groups the most common depth score is 3, defined as new legislation or a subsequent case that superseded both the point of law and the outcome of the prior decision.

Both for cases in the overridden group and for cases in the overruled group, depth of superseding event is associated with a lower shadow-precedent score. More relevant for the theory of shadow precedent, in each depth category, the shadow-precedent score is lower for cases in the overruled group compared with the overridden group. For example, when depth is 3, the mean shadow-precedent score of the overruled group is .48 lower than the mean shadow-precedent score of the overridden group ($= .91 - .43$); this difference is statistically significant.

Table 2 also reports the mean shadow-precedent score sorted by consciousness (for legislative overrides) and by explicitness (for judicial overrulings). The explicitness of the superseding event is distinct from the

13. The number of years that should be excluded may vary in different statutory contexts. As explained in Section 4.3, we test the robustness of our findings by excluding different numbers of years and then reestimating the basic model.

14. The online appendix describes the depth classification used for cases in the overruled group.

Table 2. Shadow Precedent Scores

	Overridden		Overruled		Overridden – Overruled
	Count	Mean	Count	Mean	
Full sample	132	.89	49	.58	–.31** (.001)
Depth (override or overruling):					
1	14	1.20	7	1.10	–.10 (.717)
2	16	.87	11	.65	–.22 (.277)
3	76	.91	18	.43	–.48** (.004)
4+	26	.68	13	.44	–.24+ (.083)
Consciousness of override:					
Conscious	52	.84			
Not conscious	80	.92			
Explicitness of overruling:					
Explicit			32	.48	
Not explicit			17	.77	
Type of override:					
Restorative	29	.70			
Updating or clarifying	103	.94			

Note. Values in parentheses are *p*-values for the difference in means.

+ Significant at the 10 percent level (two-sided test).

** Significant at the 1 percent level (two-sided test).

depth of the event; some cases include clear statements that a minor point in a prior decision is superseded while simultaneously affirming that the primary rule from the prior case remains good law. There is a relatively large difference between cases subject to an explicit judicial overruling (.49) and those in which the overruling is not explicit (.77); and overrides that strongly repudiate precedents (that is, a restorative override) have a far lower shadow-precedent score (mean = .70) than updating or clarifying overrides (mean = .94). By contrast, conscious legislative overrides, as defined by Buatti and Hasen (2015), have only a slightly lower shadow-precedent score (.84) than nonconscious overrides (.92).

To illustrate the effect of an override or an overruling, as compared with each other and with our control group, we track Citation Ratio_{*t*}, defined as follows:

Citation Ratio_{*t*}

$$= \frac{\text{Net citations in year } t}{\text{Average net citations per year in the pre-event period}}.$$

Cases in all three groups are subject to precedent depreciation. Thus, whatever postevent difference we observe in shadow precedent among

Figure 1. Median citation ratios by treatment group

the overridden group, the overruled group, and the matched control group can more naturally be attributed to the difference in treatment rather than simply depreciation over time.

Figure 1 shows the median citation ratio for cases in the overridden group compared with the overruled group and the control group over the 16-year event window. For most of the postevent period, net citations of the overruled group are about 20 percentage points lower than those of the other two groups. This gap persists (and indeed widens) over the full 10-year postevent period. The legislatively overridden group is almost indistinguishable from the control group.

To clarify these results, Table 3 reports for each event year the average ratio of total, positive, and negative citations divided by the average number of total citations to each case during the pre-event period ($t < 0$). Means are winsorized at 99 percent to reduce the impact of outliers. All three groups receive a small ratio (approximately 4 percent) of negative citations in the years prior to the event. After the event, the two superseded groups receive an increase in negative citations, while the control group's citations remain largely unchanged. For cases in the judicially overruled group, this is a large increase, with the average overruled case receiving 12–17 percent negative citations per year immediately following

Table 3. Average Number and Percentage of Citations per Event Year by Shepard's Category

<i>t</i>	Overridden				Overruled				Matched Control			
	<i>N</i>	Total (%)	Positive (%)	Negative (%)	<i>N</i>	Total (%)	Positive (%)	Negative (%)	<i>N</i>	Total (%)	Positive (%)	Negative (%)
-5	61	101	8	4	50	88	5	4	61	101	6	4
-4	73	104	9	4	52	89	7	5	69	95	7	4
-3	87	102	7	4	52	92	7	4	75	105	7	3
-2	99	98	9	4	52	105	13	6	83	104	7	4
-1	115	104	10	4	53	106	10	6	98	98	7	4
0	138	99	10	4	53	111	6	17	111	96	9	3
1	138	102	9	7	53	89	6	13	111	95	7	4
2	137	95	8	6	53	85	5	12	111	93	8	3
3	134	87	8	5	50	74	5	8	110	88	8	2
4	130	87	8	4	47	71	3	9	105	85	7	2
5	126	89	8	5	48	62	4	7	102	91	9	3
6	126	89	8	4	45	64	4	4	102	91	8	2
7	122	88	7	4	42	58	4	5	100	88	8	2
8	117	83	7	4	41	59	4	5	96	90	9	2
9	115	86	9	4	39	56	4	4	94	91	9	3
10	114	99	10	4	37	57	4	5	92	84	8	2

Note. Values are mean percentages (winsorized at 99 percent) of the number of total, positive, and negative annual citations to each case divided by the average number of total citations to the case during the pre-event period.

the overruling; by contrast, an average case in the overridden group never receives more than 7 percent negative citations at any point over the 10 years we track.¹⁵

Looking at total and positive citations also reveals some important differences. In the overruled group, there is a rapid postevent decline in total and positive citations. By about 5 years after the event, the precedential value of the average overruled decision is cut by almost half. In the overridden group, by contrast, the numbers of total and positive citations decline only slightly throughout our event window. Overall, the overridden group appears more like the control group than the overruled group.

4. TESTING EXPLANATIONS FOR RELIANCE ON SHADOW PRECEDENT

The results above provide tentative support for the theory of shadow precedent, emphasizing that citations decline dramatically after a judicial overruling but only minimally after a legislative override. In this section we test the hypothesis using multivariate regression analysis.

There are some limitations with using net citations as a proxy for the precedential value of the underlying decision. First, as illustrated by Table 3, for most cases the number of negative citations is dwarfed by the large number of neutral citations (and, to a lesser extent, positive citations). Yet negative citations—especially following an override or overruling—are conceptually important, as they demonstrate an acknowledged and considered reduction in the precedential value attached to the original decision. Accordingly, we believe that they should be given more weight than a string citation with no discussion of the cited case. Second, across different types of cases, there is wide disparity in the average number of annual citations and year-to-year variance in such citations. To illustrate, some of the habeas corpus cases in our overridden sample receive more than 1,000 citations per year, while the median case in the overridden group receives about 26 total citations per year. The citation pattern of cases that receive more annual citations is not inherently more important for purposes of understanding shadow precedent. Yet if we were to use

15. Interestingly, even though overrides are generally not retroactive, the highest percentage of negative citations occurs very quickly after the override is enacted. During this period, courts may be (properly) resolving decisions according to the former statutory language but also flagging that an override has been enacted that will (subsequently) change the standard applied. Such statements may be coded as negative citations.

net citations as our measure of precedent in the analysis below, such cases would be given disproportionate weight.

To address both concerns, we replace net citations with net logged citations, defined for each case (i) and each event year (t) as

$$\begin{aligned} \text{Net Logged Citations}_i = & \text{Log}(\text{Positive} + \text{Neutral} + \text{Cited By} + 1) \\ & - \text{Log}(\text{Warning} + \text{Caution} + \text{Questioned} + 1). \end{aligned}$$

Because the log function is concave and most cases receive considerably fewer negative citations than positive or neutral citations, our measure of Net Logged Citations will naturally give more weight to a small fluctuation in negative citations and less weight to modest fluctuations in positive or neutral citations to a heavily cited case. We use Net Logged Citations as our dependent variable in the empirical analysis below. In Section 4.3 we consider the robustness of our results to alternative specifications of the dependent variable.

One concern for our regression analysis is that legislative overrides and judicial overrulings are not exogenous events. The underlying political conditions and external developments that led to an override or overruling may also impact how a superseded case would have been cited even in the absence of such an event. To address this concern, we employ two identification strategies in the regression analysis below: case-level matching and headnote-level analysis.

4.1. Case-Level Regression Analysis

Using fixed-effects regression analysis, we estimate the following equation:

$$\begin{aligned} \text{Net Logged Citations}_{it} = & \alpha + \beta_1 \times \text{Post}_t + \beta_2 \times (\text{Post} \times \text{Override}_{it}) \\ & + \beta_3 \times (\text{Post} \times \text{Overrule}_{it}) \\ & + \beta \times \mathbf{X} + \text{CaseFE}_i + \varepsilon, \end{aligned} \quad (1)$$

where i indexes cases from our sample, t indexes year relative to the event $t \in [-5, 10]$, CaseFE_i is a set of case fixed effects, $\mathbf{X}$ is a vector of included control variables, and ε is the error term. Fixed-effects analysis is particularly appropriate here as it creates a pre-event baseline for each case and then compares how cases in each treatment group were cited pre-event and postevent. Fixed-effects analysis also reduces risk of omitted-variable bias by eliminating all time-constant effects, both observed and unobserved (Wooldridge 2002).

For purposes of estimating equation (1), we again exclude observations from years $t = 0, 1$, and 2 ; from the year that the case was decided; and from any cases that received fewer than three citations per year in the pre-event period. Putting these restrictions in place, we estimate equation (1) on panel data from all three groups of cases covering 2,510 years of citation data to 224 cases.

The primary explanatory variables of interest for testing the shadow-precedent hypothesis are Post_i , which equals one if $t > 0$ and zero otherwise; $\text{Post} \times \text{Override}_i$, which equals one if $t > 0$ and case i is in the legislative overridden sample and zero otherwise; and $\text{Post} \times \text{Overrule}_i$, which equals one if $t > 0$ and case i is in the judicial overruled sample and zero otherwise. We separately estimate the effect of an override as opposed to an overruling, with both coefficient estimates net of any change in citations to the control group. The shadow-precedent hypothesis predicts that $\beta_{\text{Override}} < 0$, $\beta_{\text{Overruling}} < 0$, and $\beta_{\text{Override}} > \beta_{\text{Overruling}}$, namely, that shadow-precedent scores will be higher for overrides than for overrulings.

Table 4 presents regression results, reporting fixed-effects coefficient estimates. Models 1–3 apply to the full sample, while models 4–6 include only cases in which the CEM algorithm found a 1:1 match. In addition to our primary explanatory variables, all models include year dummy variables and the variable $\text{Log}(\text{Years Since Decision})$, which reflects non-linear depreciation of precedent over time, as found in Black and Spriggs (2013).

Models 2 and 3 add explanatory variables that may help clarify to a lower-court judge the extent to which the precedent is superseded. The first such variable is $\text{Depth} \times \text{Post}$, which equals the interaction between Post and the depth score assigned to the superseding event, recoded to a 0–4 scale. We expect that deeper overrides and deeper overrulings will be associated with a lower shadow-precedent score. The second variable added is $\text{Restorative} \times \text{Post}$, which equals one if $t > 0$ and the override is classified as restorative by Christianson and Eskridge (2014) and zero otherwise. As described above, we expect restorative overrides to be associated with lower shadow-precedent scores, although the ideological division between the Court and Congress could suggest the opposite result.

We also control for subsequent Supreme Court citations to an overruled or overridden precedent. Citations by the Supreme Court provide interpretive guidance regarding the validity of the precedent and are one of the few factors that have been shown to affect the typical rate at which

precedents depreciate (Black and Spriggs 2013). We add two variables, SC Postevent Nonwarning Cites and SC Postevent Warning Cites, which equal the number of times, as of year t , that the Court has cited the original case—in either a nonwarning or a warning respect—since the event.

We also control for lower-court opinions issued shortly after the event that provide either a positive or a negative interpretation of the override or overruling. We hypothesize that such early decisions may set a path that other courts follow, even courts for which the early citation has no binding authority (Corley, Collins, and Calvin 2011). To operationalize this, we set *Sideways Cites Negative* _{it} equal to the ratio of negative citations that case i received in years 1–3 to the total number of citations that case i received over the same time period. We define *Sideways Cites Positive* _{it} similarly, except it is based on positive citations over the 3 years immediately following the event. Both of these variables are set to zero if $t \leq 0$.

Finally, in model 3 we add variables to control for whether the decision was subject to an explicit judicial overruling (*Explicit* $\times$ *Post*) or a conscious legislative override (*Conscious* $\times$ *Post*). These variables are designed to capture whether the superseding body—Congress or the Court—stated clearly that it was superseding (at least in part) a prior decision.¹⁶ For judicial overrulings, *Explicit* $\times$ *Post* equals one if $t > 0$ and the overruling decision was coded as explicit and zero otherwise. For legislative overrides, *Conscious* $\times$ *Post* equals one if $t > 0$ and the override is included in Buatti and Hasen (2015), and zero otherwise. We expect that both *Explicit* $\times$ *Post* and *Conscious* $\times$ *Post* will have a negative impact on net citations.

One of the benefits of using CEM is that it can reduce covariate imbalance between the treatment and control groups. However, to take advantage of this feature, we need to remove unmatched cases. To implement this, we reestimate models 1–3 limited to the sample of 141 cases with a 1:1 match in the control group. Again we exclude observations from years $t = 0, 1$, and 2; from the year that the case was decided; and from any cases that received fewer than three citations per year in the pre-event period. Putting these restrictions in place, we reestimate models 1–3 on panel data covering 1,874 years of citation data and 169 matched cases. Results are reported in models 4–6.

16. As explained more fully in the online appendix, they are not wholly comparable, since *Conscious* considers statements in both statutory language and key legislative history and *Explicit* looks only at the text of the controlling judicial opinion.

Table 4. Baseline Case-Level Regression Analysis

	Full Sample (N = 2,510)			Matched Cases (N = 1,873)		
	(1)	(2)	(3)	(4)	(5)	(6)
Post	-.249* (.101)	-.298* (.135)	-.280* (.141)	-.213* (.102)	-.210+ (.117)	-.185 (.122)
Post × Override	-.086 (.105)	.294+ (.164)	.278+ (.165)	-.156 (.119)	.206 (.206)	.226 (.209)
Post × Overrule	-.688** (.155)	-.372+ (.204)	-.136 (.188)	-.663** (.200)	-.295 (.238)	-.111 (.220)
Depth × Post		-.197** (.067)	-.171** (.066)		-.190* (.094)	-.175+ (.090)
Restorative × Post		-.476+ (.269)	-.468+ (.269)		-.479 (.354)	-.438 (.340)
Conscious × Post			-.103 (.187)			-.173 (.224)
Explicit × Post			-.457* (.183)			-.400+ (.211)
Sideways Cites Negative		-.648 (.638)	-.412 (.638)		-2.792** (.852)	-2.423** (.861)
Sideways Cites Positive		1.043	.899		2.081*	1.820*

SC Postevent Nonwarning Cites		(1.060)	(1.052)	(.893)	(.854)
		.026 ⁺	.020 ⁺	.028	.021
SC Postevent Warning Cites		(.012)	(.011)	(.021)	(.022)
		.296	.263	.459*	.458*
Log(Years Since Decision)		(.264)	(.259)	(.220)	(.230)
	-.056	-.040	-.040	-.052	-.053
	(.067)	(.058)	(.060)	(.064)	(.065)
Case clusters	224	224	224	169	169
R^2 (within)	.259	.299	.306	.321	.327
Wald F -tests:					
$\beta_{\text{Override}} = \beta_{\text{Override}}$	16.09**	20.67**	8.38**	5.77*	4.27*
$\beta_{\text{Override}} + \beta_{\text{Restorative}} = \beta_{\text{Override}}$		.43	.04	.00	.08
$\beta_{\text{Override}} + \beta_{\text{Conscious}} = \beta_{\text{Override}} + \beta_{\text{Explicit}}$			8.76**		3.49 ⁺

Note. The unit of analysis is annual citations to each case over a 16-year event window. The dependent variable is Net Logged Citations. To address serial correlation (Bertrand, Duflo, and Mullainathan 2004), standard errors (in parentheses) are clustered at the case level. All regressions include year dummies and case fixed effects.

⁺ Significant at the 10 percent level (two-sided test).

* Significant at the 5 percent level (two-sided test).

** Significant at the 1 percent level (two-sided test).

In each model reported in Table 4, we find results consistent with the shadow-precedent hypothesis. Judicial overrulings have a stronger negative effect on postevent citations than do legislative overrides. The coefficient estimates for $\text{Post} \times \text{Override}$ ($-.16$ to $.29$) and $\text{Post} \times \text{Overrule}$ ($-.69$ to $-.11$) make this clear. Indeed, in each model, the estimate for β_{Overrule} is at least $.33$ less than the estimate for β_{Override} , and in each model we can confidently reject, using a Wald F -test, the null hypothesis that $\beta_{\text{Override}} = \beta_{\text{Overrule}}$. By contrast, the coefficient estimate for β_{Override} is significant only in two of the six models reported in Table 4, which means that citations to the overridden group are not significantly different from citations to the control group in all models.

As expected, we find that, in all models, $\text{Depth} \times \text{Post}$ has a significant negative effect on postevent citations. Nonetheless, after controlling for depth, we still find a significantly larger drop in the number of citations to cases overruled by a judicial opinion than cases overridden by Congress. Model 3 also shows that cases subject to an explicit overruling or a conscious override receive fewer citations (significantly so for explicit overrulings) after the event but that overruled cases have lower shadow-precedent scores than overridden cases (that is, $\beta_{\text{Override}} + \beta_{\text{Conscious}} > \beta_{\text{Overrule}} + \beta_{\text{Explicit}}$). These findings demonstrate that our results are not driven by comparing high-depth explicit overrulings with low-depth unconscious overrides, but rather in each category we find more reliance on shadow precedent following a legislative override than following a judicial overruling.

Finally, we compare restorative overrides with judicial overrulings. Table 4 shows that, everything else (including depth) being equal, we cannot reject that $\beta_{\text{Override}} + \beta_{\text{Restorative}} = \beta_{\text{Overrule}}$. In other words, while shallow and nonrestorative overrides have little effect on postevent citations, the relatively small group of cases superseded by a restorative override experience a decline in net citations that is similar to that for cases subject to judicial overrulings of comparable depth. Since restorative overrides highlight an interpretive or ideological disagreement between Congress and the Court, the fact that citing courts are more responsive to restorative overrides than to updating or clarifying overrides (even those that are similarly deep) suggests that shadow precedent is more likely caused by information failure than by lower courts seeking to align themselves with the Supreme Court in areas of dispute between the Court and Congress.

4.2. Headnote-Level Regression Analysis

Our measurement of shadow precedent in the analysis above is clouded by the fact that a case may stand for several legal propositions, only some of which were impacted by the superseding event. Consequently, some of the citations made after an override or overruling are presumably to unrelated legal propositions, which adds noise to our empirical analysis. Our variable *Depth* captures this to some extent, as it distinguishes between cases in which the core reasoning is fully repudiated by an override or overruling and those that are only minimally affected. In this section, we develop an additional novel identification strategy that explicitly addresses the fact that each case represents multiple legal propositions.

LexisNexus uses distinct headnotes to divide cases into separate legal propositions and then tracks citations to each headnote in a case. Taking advantage of this feature, we randomly selected 60 cases from the overridden sample and 20 cases from the overruled sample (in each group, this represents approximately 36 percent of the sample). We then hand coded each headnote for the cases in this subsample, using the following three classifications:¹⁷ directly superseded by the new statute or overruling case (category 1), arguably superseded by the new statute or overruling case (category 2), or unrelated to the new statute or overruling case (category 3). This effectively gives us multiple levels of treatment, and we can compare how directly superseded (category 1) propositions are cited after an override or overruling with arguably superseded (category 2) or unrelated (category 3) propositions. An advantage of this approach is that all of the observations come from exactly the same fact pattern. Consequently, unobserved features of each case, even time-varying features, are unlikely to be a source of bias because they apply to all three categories.

Table 5 reports the average (mean) ratio of total and negative annual citations divided by the average number of total citations to each headnote during the pre-event period. Pre-event, each headnote category receives a small ratio of negative citations. Postevent, there is a meaningful increase in negative citations to category 1 headnotes, particularly for the judicially overruled group. For example, negative citations to category 1 headnotes in the overruled group increase from approximately 7 percent per year pre-event to 19 percent to 31 percent in the 5 years immediately

17. Because headnote coding is complex and labor intensive, we did not classify headnotes for all the cases in our full sample. We performed an intercoder reliability check; 74 percent of the headnotes were classified identically. For more details on this process, see the online appendix.

Table 5. Average Number and Percentage of Citations per Year by Headnote and Shepard's Category

<i>t</i>	Directly Superseded: Category 1			Arguably Superseded: Category 2			Completely Unrelated: Category 3		
	<i>N</i>	Total (%)	Negative (%)	<i>N</i>	Total (%)	Negative (%)	<i>N</i>	Total (%)	Negative (%)
Legislative overrides:									
-5	31	81	3	67	87	3	10	78	1
-4	38	99	6	82	96	6	36	108	8
-3	50	97	8	90	94	4	42	91	5
-2	57	99	6	106	93	3	42	91	8
-1	69	94	6	117	107	6	44	103	8
0	69	88	7	117	96	6	44	97	8
1	69	89	10	117	104	7	44	89	7
2	69	67	6	117	85	4	44	89	6
3	69	71	7	117	91	6	44	78	4
4	65	61	2	115	87	3	42	86	1
5	62	70	5	111	81	3	38	85	2
6	62	73	2	111	85	2	38	84	5
7	58	81	4	111	99	4	38	89	3
8	57	81	3	105	101	4	30	79	3
9	57	104	4	103	125	4	30	92	3
10	57	117	4	103	146	6	30	127	4
Mean pre-event	49	94	6	92.4	95	4	34.8	94	6
Mean postevent	62.5	81	5	111	100	4	37.8	90	4

Judicial overrulings:									
-5	27	87	7	43	68	5	11	98	7
-4	27	91	7	43	79	9	11	109	11
-3	27	100	11	43	102	12	11	119	3
-2	27	106	5	43	124	8	11	78	1
-1	27	98	6	43	113	5	11	86	5
0	27	117	31	43	114	24	11	110	15
1	27	107	24	43	100	16	11	104	8
2	27	99	21	43	91	14	11	78	9
3	27	82	21	43	74	12	11	105	11
4	27	79	19	43	84	10	11	97	6
5	27	52	19	43	59	10	11	85	9
6	26	55	8	38	75	8	11	100	5
7	26	57	11	38	73	7	11	84	7
8	25	58	7	37	79	9	11	128	2
9	25	39	10	37	57	8	11	103	4
10	25	45	13	37	66	12	11	73	4
Mean pre-event	27	97	7	43	97	8	11	98	5
Mean postevent	26.2	67	15	40.2	76	11	11	96	7

Note. Values are mean percentages (winsorized at 99 percent) of the number of total and negative annual citations to each headnote divided by the average number of annual citations to the headnote during the pre-event period.

following the overruling. We find far less change in citation patterns to category 2 and 3 headnotes.

To test the shadow-precedent hypothesis on the headnote-level data, we use fixed-effects regression analysis to estimate the following equation:

$$\begin{aligned} \text{Net Logged Citations}_{th} = & \alpha + \beta_1 \times \text{Post}_t + \beta_2 \times (\text{Post} \times \text{Category1}_{th}) \\ & + \beta_3 \times (\text{Post} \times \text{Category2}_{th}) \\ & + \beta \times \mathbf{X} + \text{HeadnoteFE}_h + \varepsilon, \end{aligned}$$

where h indexes individual headnotes from the subsample cases, t indexes the year relative to the event $t \in [-5, 10]$, HeadnoteFE_h is a fixed effect for each headnote, $\mathbf{X}$ is a vector of included control variables, and ε is the error term. In the headnote context, the dependent variable is defined as

$$\begin{aligned} \text{Net Logged Citations}_{th} = & \text{Log}(\text{Total} - \text{Negative} + 1) \\ & - \text{Log}(\text{Negative} + 1). \end{aligned}$$

We exclude observations from years $t = 0, 1$, and 2 ; from the year that the case was decided; and from any headnote that received fewer than two citations per year in the pre-event period.¹⁸ With these restrictions in place, we estimate equation (2) on panel data covering 3,453 years of headnote-level citation data from a group of 330 headnotes.

In the headnote context, the primary explanatory variables of interest are Post_t , $\text{Post} \times \text{Category1}_{th}$, and $\text{Post} \times \text{Category2}_{th}$. The interaction terms measure the marginal difference in postevent net logged citations between categories 1, 2, and 3. Similar to the above process, we predict that $\beta_{\text{Category1}} < 0$, $\beta_{\text{Category2}} < 0$, and $\beta_{\text{Category1}} < \beta_{\text{Category2}}$. Table 6 presents these results. Model 7 shows that postevent citations are significantly lower for category 1 headnotes relative to category 3 headnotes ($\beta_{\text{Category1}} = -.75$), but category 2 headnotes are not significantly different from category 3 headnotes. This suggests that subsequent citations to portions of an opinion that provide reasoning or background related to a proposition that was overridden or overruled (but that are not themselves directly superseded) are little affected by the event. Table 6 also shows that we can confidently reject the null hypothesis that $\beta_{\text{Category1}} = \beta_{\text{Category2}}$.

To explore whether there is a difference in citation patterns to su-

18. Since headnotes receive, on average, many fewer citations than the case as a whole, we use two citations per year in the pre-event period as a minimum threshold for inclusion in this analysis rather than the three citations per year minimum that we used in the case-level analysis.

Table 6. Headnote-Level Regression Analysis

	(7)	(8)
Post	-.263** (.100)	-.268* (.109)
Post × Category1	-.745** (.138)	-.562** (.151)
Post × Category2	-.138 (.107)	-.073 (.130)
Post × Overruled		-.135 (.195)
Post × Category1 × Overruled		-.403 (.298)
Post × Category2 × Overruled		-.161 (.234)
R ² (within)	.169	.178
Wald F-test:		
$\beta_{\text{Category1}} = \beta_{\text{Category2}}$	22.13**	10.58**
$\beta_{\text{Overruled}} + \beta_{\text{Category1}} \times \text{Overruled} = 0$		5.40 ⁺

Note. Values are fixed-effects regression estimates on annual citations to headnotes from decisions that were subject to legislative override ($n = 60$) or judicial overruling ($n = 20$) between 1985 and 2011. The unit of analysis is annual citations to each headnote over a 16-year period surrounding the event. The dependent variable is Net Logged Citations. To address serial correlation (Bertrand, Duflo, and Mullainathan 2004), standard errors (in parentheses) are clustered at the headnote level. All regressions include year dummies and headnote fixed effects. $N = 3,453$; headnote clusters = 330.

⁺ Significant at the 10 percent level (two-sided test).

* Significant at the 5 percent level (two-sided test).

** Significant at the 1 percent level (two-sided test).

perseded headnotes after a legislative override as opposed to a judicial overruling, we interact the explanatory variables used in model 7 with Overruled. This yields three new variables: Post × Overruled, Post × Category1 × Overruled, and Post × Category2 × Overruled. Adding these variables to model 8, we find that, consistent with the shadow-precedent hypothesis, there is a larger drop in the number of citations to category 1 headnotes following a judicial overruling compared with a legislative override: we reject the null hypothesis that $\beta_{\text{Overruled}} + \beta_{\text{Category1}} \times \text{Overruled} = 0$.

4.3. Robustness Checks and Alternative Explanations

In this section we explore the robustness of our results to alternative measurements of precedent, alternative event windows, and the use of a balanced versus unbalanced panel. We also discuss our efforts to assess the role that ideological differences between the courts and Congress may play.

4.3.1. *Alternative Measurements of Precedent.* Some studies that use citation data exclude neutral and cited-by references from consideration on the theory that such citations reveal little about the prudential significance of a case (for example, Westerland et al. 2010).¹⁹ Our study includes neutral and cited-by references in the positive category because we felt that a neutral discussion of a case that does not flag the fact of an override or overruling—which would have resulted in a negative warning—is, for our purposes, a positive citation in the sense that it treats the precedent as presumptively valid. To investigate whether our results depend on this choice, we create an alternative dependent variable based only on positive and negative citations, with other citation categories—namely, neutral and cited by—removed from the analysis. Our modified dependent variable is set equal to $\log(\text{Positive} + 1) - \log(\text{Negative} + 1)$. We then reestimate model 2 using this alternative dependent variable. Our results, reported in model 9 of Table 7, are qualitatively unchanged. We still find that there is a greater decline after a judicial overruling than after an override, and early interpretive guidance operates similarly to the models reported in Table 4.

4.3.2. *Nonretroactivity of Legislative Overrides versus Retroactivity of Judicial Overrulings.* Judicial overrulings typically are retroactive, whereas legislative overrides typically are not. Thus, judicial overrulings usually take effect immediately, and all cases decided after the date of the overruling ($t = 0$) should be decided under the new standard. By contrast, cases decided shortly after an override will usually be based on the old statutory language, and thus it may still be appropriate to cite the pre-existing precedent interpreting the prior statutory language. To address

19. There are also some studies that use the total number of citations without distinguishing between positive, neutral, and negative citations (for example, Fowler et al. 2007; Cross et al. 2010). The structure of our study depends on the distinctions between positive and negative citations, and negative citations are far more prevalent in our data set, which consists of overridden and overruled cases, than in most other studies. Accordingly, grouping all citations together was not a viable approach for our study.

this problem, we use an asymmetric event window, starting 5 years before the event and running until 10 years after it. In addition, in the analysis above, we exclude the data from years $t = 0, 1$, and 2 when we expect that the nonretroactivity problem is most likely. The choice to exclude 3 years of data, however, is admittedly arbitrary. There are presumably some cases decided after this 3-year period that are properly resolved according to the standard that predated the override, and there are surely some cases decided during the 3-year period that should be governed by the amended statute.

To investigate whether our results are sensitive to the choice of how many years to exclude, we test two alternative approaches. First, we reestimate model 2, excluding 6 years ($t = 0$ to $t = 5$) of observations. Results are reported in model 10 in Table 7. Though we lose almost 700 observations by expanding the nonretroactivity period, this does not qualitatively change our main findings. Second, we reestimate model 2, excluding just the year of the event ($t = 0$). Results are reported in model 11. Again, our results remain largely consistent. Collectively, these models suggest that our results are not driven by delayed application of legislative overrides.

4.3.3. Unbalanced versus Balanced Panel Data. Our study uses unbalanced panel data; the superseding event sometimes occurs fewer than 10 years prior to 2013, the last year for which we collected citation data (effectively truncating the postevent period) or fewer than 5 years after the original decision (truncating the pre-event period). If the missing observations due to truncation or gaps in the data were random, the use of an unbalanced panel would not cause a concern. In our case, however, observations at the start of our pre-event window likely reflect a nonrandom subsample of cases, because restorative overrides tend to be enacted much more quickly than nonrestorative overrides (and thus are disproportionately likely to be excluded from that portion of the pre-event period).

To address this concern, we shorten the event period to a single pre-event observation ($t = -1$) and a single postevent observation ($t = 3$) for each case (that is, the pre- and postevent observations that are closest in time to the event, other than the years excluded because of the nonretroactivity issue). This yields a balanced sample, with exactly two observations for each case. We then reestimate model 2, with results reported in model 12 (Table 7). We find less reliance on shadow precedent following

Table 7. Robustness Checks

	(9)	(10)	(11)	(12)	(13)
Post	-.040 (.091)	-.410* (.184)	-.054 (.113)	-.073 (.156)	-.331* (.165)
Post $\times$ Override	.004 (.161)	.264 (.163)	.210 (.147)	.089 (.185)	.291+ (.163)
Post $\times$ Overrule	-.477** (.177)	-.312 (.199)	-.360* (.178)	-.956** (.343)	-.341 (.228)
Depth $\times$ Post	-.096 (.071)	-.216** (.065)	-.155* (.061)	-.156+ (.082)	-.194** (.066)
Restorative $\times$ Post	-.432* (.211)	-.327 (.303)	-.267 (.225)	-.654* (.306)	-.462+ (.272)
Sideways Cites Negative	-1.035** (.313)	-.736 (.623)	-.466 (.660)	-.475 (.689)	-.659 (.638)
Sideways Cites Positive	.283 (.557)	1.425 (1.109)	.687 (.983)	-.562 (1.069)	1.051 (1.060)
SC Postevent Nonwarning Cites	.022+ (.013)	.029* (.012)	.021+ (.012)	.194** (.049)	.026** (.011)
SC Postevent Warning Cites	.289 (.235)	.206 (.260)	-.038 (.213)	.130 (.442)	.284 (.262)

a judicial overruling, and we can confidently reject the null hypothesis (that is, $\beta_{\text{Override}} \neq \beta_{\text{Overrule}}$). We prefer the regression analysis in Table 4, using unbalanced panel data, because the longer postevent window can better detect the full effect of the superseding event and the longer pre-event window is less susceptible to an unusual citation pattern in the year immediately preceding the override or overruling. Nonetheless, it is reassuring that our choice to use an unbalanced panel does not seem to be driving our results.

4.3.4. Ideology. Positive political science models typically understand overrides as constraints on the Supreme Court's ability to interpret statutes in line with its ideological preferences. We hypothesized, relatedly, that lower courts might feel pressure to interpret overrides narrowly to conform to the Supreme Court's presumed preferences. Lower-court judges might also use the ambiguity implicit in determining how overrides relate to the precedent to advance their own ideological preferences. Our data do not permit us to measure the ideology of individual trial court judges citing the overridden precedent. However, we assess the ideological direction of the override in relation to which party controlled Congress and the presidency at the time of the later decisions. This provides a rough proxy for the likelihood that an unreasonably narrow interpretation of the override (to conform with the presumed preference of the Supreme Court or the lower court's own ideological preferences) would trigger a second override.

In model 13, we reestimate model 2 with three new explanatory variables: Fed Aligned with Override, equals 0–3 on the basis of whether in year t the president, the House of Representatives, and/or the Senate is from the same party as the direction that the override moved the law relative to the precedent.²⁰ So, for example, if the override moved the law in a liberal direction relative to the precedent and if the president and a majority of the Senate (but not a majority of the House) were from the Democratic party, we would set Fed Aligned with Override, equal to 2 in year t . We also include an explanatory variable Unified Control of Government, which equals one when one political party controls the presidency and both houses of Congress and zero otherwise, as Congress may be less likely to pass new legislation in periods of divided government

20. For each override, the classification of the direction in which it moved the law was provided to us by Matt Christianson and Bill Eskridge in conjunction with their study (Christianson and Eskridge 2014).

(see, for example, Coleman 1999). We also control for whether the original case was a liberal decision by adding the variable *Liberal Case* $\times$ *Post*. We find no evidence that reliance on shadow precedent is driven by ideology. Also, as discussed above, restorative overrides tend to have lower shadow-precedent scores than nonrestorative overrides, even though the ideological divides between Congress and the Court are far more pronounced in this context. However, our measure of ideology is admittedly crude, and this may be a fruitful area for further research.

5. IMPLICATIONS AND CONCLUSION

The assumption that a legislative override is equivalent to a judicial overruling is incorrect. After a judicial overruling, the number of total citations and net citations drops quickly and sharply, and negative citations become quite common. Within about 6 years after the event, the precedential weight of the decision has been cut roughly in half. But after a congressional override, the number of total citations and net citations drops only a little, and negative citations remain relatively rare. Many overridden decisions are still widely cited even 10 years after the override. In addition, although there is a more noticeable decrease in citations after a relatively deep override than after a shallow override, at every level of depth, the number of citations drops more rapidly after a judicial overruling than after a statutory override. Existing debate has centered on the extent to which court action is constrained by the threat of a potential override. Our findings suggest that, even after an override, courts may be unconstrained.

Our data do not establish definitively what causes the differences we observe, but we suggest that information failure or judicial error are important factors. Litigants and courts simply may not realize that a statutory provision has been enacted that calls into question the validity of the precedent. Consistent with this, we found that restorative overrides—which are more likely to address specific judicial precedents in statutory language or legislative history and are more likely to be heavily publicized—result in lower levels of shadow precedent than nonrestorative overrides. We also hypothesized that lower courts, even if aware of an override, might be unsure how to synthesize it with precedent. We predicted that clear signals from Congress and from other judicial actors could reduce reliance on overridden precedents, and we find results

broadly consistent with this explanation. It is also possible that courts use the ambiguity implicit in overrides to advance ideological preferences, either their own or the presumed preferences of reviewing courts. Although our results do not establish this, future researchers may wish to design tests to assess the potential impact of ideology on interpretation of overrides more directly.

We were particularly struck by our findings regarding nonrestorative overrides, those that are intended to update or clarify statutory law. Christiansen and Eskridge (2014) demonstrate that the Supreme Court often explicitly invites such overrides on the grounds that they should be enacted by the legislative branch rather than implemented through statutory interpretation by unelected judges. Congress heeds these calls by enacting new statutory provisions. But we find that, in many respects, citation levels to cases overridden by such statutes are very similar to citation levels to cases in our control group.

One would hope, of course, that lawyers would bring all relevant statutory developments to the attention of courts and that courts, in any case, would properly apply the controlling statutory law. To probe this question further—and to help distinguish information failures from other potential explanations—it would be helpful to analyze whether lawyers' briefing regarding overrides affects courts' ongoing reliance on shadow precedents. That analysis is beyond the scope of our project, but perhaps it can be explored in future research. Even in the absence of empirical evidence on point, it seems apparent that lawyers should more fully integrate the analysis of overrides in crafting their legal arguments. Lawyers need to carefully read the statutory language that governs resolution of a dispute and consider whether judicial decisions interpreting the statute predate any changes to the statutory language. This is true not only for Supreme Court decisions but also for lower-court decisions that may rely on Supreme Court precedent that has been superseded. Courts, likewise, are charged with interpreting and implementing existing law, and they may be expected to do such research even if the lawyers appearing before them have not properly briefed changes (see Widiss [2015] for a fuller discussion of the respective responsibilities of lawyers and courts).

LexisNexis, Westlaw, and other legal research resources could also reconsider the coding protocols that they use for flagging that a judicial precedent has been affected by subsequent statutory actions. As described more fully in Widiss (2014), these services rely primarily on judicial statements indicating that a statutory amendment has superseded a precedent

before flagging the precedent. This necessarily builds in a lag time, which is often several years. Relying on judicial signals makes sense in a purely common-law-based system, but the approach may be reconsidered for statutory decisions.

Even more fundamentally, if Congress seeks, in enacting an override, to end reliance on the relevant portion of the preexisting precedent, congressional drafters should make the relationship between statutory amendments and prior case law clear in the statutory language. Administrative agencies could also help explain how statutory amendments affect the validity of precedents. These changes could facilitate prompt flagging by legal research databases and make it easier for lawyers to understand the extent to which (if any) the precedent remains controlling.

Our findings suggest that overrides often fail to actually override. This is a significant problem for bedrock principles of legislative supremacy. If Congress is, in fact, to serve as the primary source of statutory law, all of these actors—Congress, administrative agencies, legal research services, lawyers, and ultimately courts—need to endeavor to ensure that overrides are implemented effectively.

REFERENCES

- Bailey, Michael A., and Forrest Maltzman. 2011. *The Constrained Court: Law, Politics, and the Decisions Justices Make*. Princeton, NJ: Princeton University Press.
- Barnes, Jeb. 2004. *Overruled? Legislative Overrides, Pluralism, and Contemporary Court-Congress Relations*. Stanford, CA: Stanford University Press.
- Benesh, Sara C., and Malia Reddick. 2002. Overruled: An Event History Analysis of Lower Court Reaction to Supreme Court Alteration of Precedent. *Journal of Politics* 64:534–50.
- Benjamin, Stuart Minor, and Georg Vanberg. 2016. Judicial Retirements and the Staying Power of U.S. Supreme Court Decisions. *Journal of Empirical Legal Studies* 13:5–26.
- Bergara, Mario, Barak Richman, and Pablo T. Spiller. 2003. Modeling Supreme Court Strategic Decision Making: The Congressional Constraint. *Legislative Studies Quarterly* 28:247–80.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. How Much Should We Trust Differences-in-Differences Estimates? *Quarterly Journal of Economics* 119:249–75.
- Black, Ryan C., and James F. Spriggs II. 2013. The Citation and Depreciation of U.S. Supreme Court Precedent. *Journal of Empirical Legal Studies* 10:325–58.

- Blackwell, Matthew, Stefano Iacus, Gary King, and Giuseppe Porro. 2009. CEM: Coarsened Exact Matching in Stata. *Stata Journal* 9:524–46.
- Boyd, Christina L., and James F. Spriggs II. 2009. An Examination of Strategic Anticipation of Appellate Court Preferences by Federal District Court Judges. *Washington University Journal of Law and Policy* 29:37–81.
- Brenner, Saul, and Harold J. Spaeth. 1995. *Stare Indecisus: The Alteration of Precedent on the Supreme Court, 1946–1992*. New York: Cambridge University Press.
- Buatti, James, and Richard L. Hasen. 2015. Conscious Congressional Overriding of the Supreme Court, Gridlock, and Partisan Politics. *Texas Law Review See Also* 93:263–88.
- Choi, Stephen J., and G. Mitu Gulati. 2008. Bias in Judicial Citations: A Window into the Behavior of Judges? *Journal of Legal Studies* 37:87–129.
- Choi, Stephen J., G. Mitu Gulati, and Eric A. Posner. 2012. What Do Federal District Judges Want? An Analysis of Publications, Citations, and Reversals. *Journal of Law, Economics, and Organization* 28:518–49.
- Christiansen, Matthew R., and William N. Eskridge, Jr. 2014. Congressional Overrides of Supreme Court Statutory Interpretation Decisions, 1967–2011. *Texas Law Review* 92:1317–1515.
- Christiansen, Matthew R., William N. Eskridge Jr., and Samuel N. Thypin-Bermeo. 2015. The Conscious Congress: How Not to Define Overrides. *Texas Law Review See Also* 93:289–315.
- Coleman, John J. 1999. Unified Government, Divided Government, and Party Responsiveness. *American Political Science Review* 93:821–35.
- Corley, Pamela C., Paul M. Collins, Jr., and Bryan Calvin. 2011. Lower Court Influence on U.S. Supreme Court Opinion Content. *Journal of Politics* 73:31–44.
- Cross, Frank B. 1997. Political Science and the New Legal Realism: A Case of Unfortunate Interdisciplinary Ignorance. *Northwestern University Law Review* 92:251–326.
- Cross, Frank B., James F. Spriggs II, Timothy R. Johnson, 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*, pp. 489–575.
- Elhauge, Einer. 2002. Preference-Estimating Statutory Default Rules. *Columbia Law Review* 102:2027–2161.
- Eskridge, William N., Jr. 1991a. Overriding Supreme Court Statutory Interpretation Decisions. *Yale Law Journal* 101:331–455.
- . 1991b. Reneging on History? Playing the Court/Congress/President Civil Rights Game. *California Law Review* 79:613–84.
- . 1994. *Dynamic Statutory Interpretation*. Cambridge, MA: Harvard University Press.
- Ferejohn, John, and Barry Weingast. 1992. Limitations of Statutes: Strategic Statutory Interpretation. *Georgetown Law Journal* 80:565–82.

- Fowler, James H., Timothy R. Johnson, James F. Spriggs II, Sangick Jeon, and Paul J. Wahlbeck. 2007. Network Analysis and the Law: Measuring the Legal Importance of Precedents at the U.S. Supreme Court. *Politics Analysis* 15:324–46.
- Gely, Rafael, and Pablo T. Spiller. 1990. A Rational Choice Theory of Supreme Court Statutory Decision with Application to the *State Farm* and *Grove City* Cases. *Journal of Law, Economics, and Organization* 6:263–300.
- Hansford, Thomas G., and James F. Spriggs II. 2006. *The Politics of Precedent on the U.S. Supreme Court*. Princeton, NJ: Princeton University Press.
- Hasen, Richard L. 2013. End of the Dialogue? Political Polarization, the Supreme Court, and Congress. *Southern California Law Review* 86:205–61.
- Hausegger, Lori, and Lawrence Baum. 1998. Behind the Scenes: The Supreme Court and Congress in Statutory Interpretation. Pp. 224–47 in *Great Theatre: The American Congress in the 1990s*, edited by Herbert F. Weisberg and Samuel C. Patterson. New York: Cambridge University Press.
- Kim, Pauline T. 2007. Lower Court Discretion. *New York University Law Review* 82:383–442.
- . 2009. Deliberation and Strategy on the United States Courts of Appeals: An Empirical Exploration of Panel Effects. *University of Pennsylvania Law Review* 157:1319–81.
- Klein, David E. 2002. *Making Law in the United States Courts of Appeals*. New York: Cambridge University Press.
- Klerman, Daniel. 2007. Jurisdictional Competition and the Evolution of the Common Law. *University of Chicago Law Review* 74:1179–1226.
- Landes, William M., and Richard A. Posner. 1976. Legal Precedent: A Theoretical and Empirical Analysis. *Journal of Law and Economics* 19:249–307.
- Levi, Edward H. 1949. *An Introduction to Legal Reasoning*. Chicago: University of Chicago Press.
- Lindquist, Stephanie A., and Frank B. Cross. 2005. Empirically Testing Dworkin's Chain Novel Theory: Studying the Path of Precedent. *New York University Law Review* 80:1156–1206.
- Luse, Jennifer K., Geoffrey McGovern, Wendy L. Martinek, and Sara C. Benesh. 2009. "Such Inferior Courts . . .": Compliance by Circuits with Jurisprudential Regimes. *American Politics Research* 37:75–106.
- Marshall, Lawrence C. 1989. "Let Congress Do It": The Case for an Absolute Rule of Statutory Stare Decisis. *Michigan Law Review* 88:177–238.
- Merryman, John Henry. 1954. The Authority of Authority: What the California Supreme Court Cited in 1950. *Stanford Law Review* 6:613–73.
- Pacelle, Richard L., Jr., Brett W. Curry, and Bryan W. Marshall. 2011. *Decision Making by the Modern Supreme Court*. New York: Cambridge University Press.
- Randazzo, Kirk A. 2008. Strategic Anticipation and the Hierarchy of Justice in U.S.

- District Courts. *American Politics Research* 36:669–93.
- Richards, Mark J., and Herbert M. Kritzer. 2002. Jurisprudential Regimes in Supreme Court Decision Making. *American Political Science Review* 96:305–20.
- Schauer, Fred. 1987. Precedent. *Stanford Law Review* 39:571–605.
- Segal, Jeffrey A. 1997. Separation-of-Powers Games in the Positive Theory of Congress and Courts. *American Political Science Review* 91:28–44.
- Segal, Jeffrey A., and Harold J. Spaeth. 2002. *The Supreme Court and the Attitudinal Model Revisited*. New York: Cambridge University Press.
- Songer, Donald R., Jeffrey A. Segal, and Charles M. Cameron. 1994. The Hierarchy of Justice: Testing a Principal-Agent Model of Supreme Court–Circuit Court Interactions. *American Journal of Political Science* 38:673–96.
- Songer, Donald R., and Reginald S. Sheehan. 1990. Supreme Court Impact on Compliance and Outcomes: *Miranda* and *New York Times* in the Courts of Appeals. *Western Political Quarterly* 43:297–316.
- Spiller, Pablo T., and Rafael Gely. 1992. Congressional Control or Judicial Independence: The Determinations of U.S. Supreme Court Labor-Relations Decisions, 1949–1988. *RAND Journal of Economics* 23:463–92.
- Staudt, Nancy C., René Lindstädt, and Jason O'Connor. 2007. Judicial Decisions as Legislation: Congressional Oversight of Supreme Court Tax Cases, 1954–2005. *New York University Law Review* 82:1340–1402.
- Sunstein, Cass R., David Schkade, Lisa M. Ellman, and Andres Sawicki. 2006. *Are Judges Political? An Empirical Analysis of the Federal Judiciary*. Washington, DC: Brookings Institution Press.
- Tokson, Matthew. 2015. Judicial Resistance and Legal Change. *University of Chicago Law Review* 82:901–73.
- Westerland, Chad, Jeffrey A. Segal, Lee Epstein, Charles M. Cameron, and Scott Comparato. 2010. Strategic Defiance and Compliance in the U.S. Courts of Appeals. *American Journal of Political Science* 54:891–905.
- Widiss, Deborah A. 2009. Shadow Precedents and the Separation of Powers: Statutory Interpretation of Congressional Overrides. *Notre Dame Law Review* 84:511–83.
- . 2012. Undermining Congressional Overrides: The Hydra Problem in Statutory Interpretation. *Texas Law Review* 90:859–942.
- . 2014. Identifying Congressional Overrides Should Not Be So Hard. *Texas Law Review* See Also 92:145–69.
- . 2015. Still Kickin' after All These Years: *Sutton* and *Toyota* as Shadow Precedents. *Drake Law Review* 63:919–46.
- Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.

ONLINE APPENDIX:

After the Override: an Empirical Analysis of Shadow Precedent

Appendix A: Identifying and Classifying Overridden and Overruled Cases

Conceptually, it is straightforward to define both statutory overrides and judicial overrulings. They are events (either newly-enacted statutory provisions or new judicial opinions) that supersede a prior judicial decision, thus changing the applicable law that will govern a given factual scenario. Often it is obvious that a given statute or decision meets this definition. At the margins, however, the line between overriding and codifying (Christiansen & Eskridge (2014), p. 1325), and the line between overruling and distinguishing (Brenner & Spaeth (1995), p. 20), may be unclear. Relatedly, both overrides and overrulings vary as to depth (i.e., how completely do they supersede the rule announced by the prior case?) and explicitness (i.e., how clear was Congress or the Court that it was replacing a rule from a prior judicial decision?).

This appendix provides additional information about how the samples of overridden and overruled cases that we use in our study were identified; how depth and explicitness were assessed; and how questions of identification and classification may affect our results.

Overridden Cases

We define our sample of overridden cases using data from Christiansen & Eskridge (2014). This study sought to identify all statutory interpretation cases that had been overridden by statutes enacted 1967-2011. It updated and expanded on an earlier study by Eskridge (1991) that had sought to identify all cases overridden by statutes enacted 1967-1990.

The Eskridge studies define an override as a statutory provision that:

- (1) Completely overrules the holding of a statutory interpretation decision, just as a subsequent Court would overrule an unsatisfactory precedent; (2) modifies the result of a decision in some material way, such that the same case would have been decided differently; or (3) modifies the consequences of the decision, such that the same case would have been decided in the same way but subsequent cases would be decided differently.

Eskridge (1991), p. 332 n.1. Christiansen & Eskridge (2014) used both legislative history and subsequent judicial indications that a statute had superseded a prior decision to identify potential overrides (Christiansen & Eskridge (2014); Widiss (2014)). Because they relied on subsequent judicial opinions, they suggest their list (and thus our list) may be incomplete for more recent years (Christiansen & Eskridge (2014), p. 1342-43). As discussed below, recent studies by Hasen (2013) and Buatti & Hasen (2015) adopt a different approach to identifying overrides that we also incorporate into our analysis.

Type of Override

Christiansen & Eskridge (2014) develops a new taxonomy for overrides, distinguishing between overrides that are “restorative,” defined as those that repudiate a prior judicial interpretation as contrary to the original Congressional intent, and those that update or clarify the law but do not indicate that the original judicial interpretation was a misinterpretation of the old law.

About one-fifth of the overrides in Christiansen & Eskridge (2014) are restorative. When passing restorative overrides, Congress is often quite clear about its frustration with the prior judicial ruling. For example, in the Lilly Ledbetter Fair Pay Act, Congress stated in the statute itself that the Court's prior decision in *Ledbetter v. Goodyear Tire & Rubber Co.*, 550 U.S. 618 (2007) had "significantly impair[ed] statutory protections against discrimination in compensation that Congress established and that have been bedrock principles of American law for decades." (Pub. L. No. 111-2 (2009)). Not all restorative overrides are this strident, and such context may be found in legislative history rather than statutory language; nonetheless, most unambiguously reject the prior precedent. Additionally, because of the power struggle inherent in restorative overrides, they often receive significant coverage in legal and popular press—and such coverage typically focuses on the "fight" between Congress and the Court. They are prevalent in areas such as civil rights or employment discrimination, where there are sharp partisan divides and where Congress in recent years has sometimes been more liberal than the Supreme Court.

Approximately two-thirds of the overrides in Christiansen & Eskridge (2014) are classified as "updating" overrides. In passing such overrides, Congress does not denigrate the prior ruling; it simply puts in place a new statutory standard that is better aligned with current policy priorities (Christiansen & Eskridge (2014), at p. 1370). Updating overrides are a mechanism for responding to changes and developments in the law or society. For example, the Digital Millennium Copyright Act of 1998, Pub. L. 105-304 (1998), created new liability for circumvention of copyright protection mechanisms. In so doing, it overrode a portion of *Sony Corp. of America v. Universal City Studios*, 464 U.S. 417 (1984), which had held that sale of VCRs that permitted copying protected content did not violate copyright laws. Notably, in *Sony*, the Supreme Court had invited such a policy intervention by Congress, stating explicitly that the Court had to be "circumspect in construing the scope of rights creative by legislative enactment" where Congress had not charted the path for how new technological developments should be handled. *Id.* at 431. Major overhauls of statutory law, such as revisions to the tax code, the bankruptcy code, civil procedure rules, intellectual property laws, or habeas statutes often include numerous updating overrides.

Finally, Christiansen & Eskridge (2014) classify about 10% of the overrides as primarily "clarifying" the law. For example, the Fraud Enforcement and Recovery Act of 2009, Pub. L. 111-21, overrode *United States v. Santos*, 553 U.S. 507 (2007), a fractured decision interpreting the federal money laundering statute. Because *Santos* did not have a majority opinion, lower courts were not sure what rule was established by the case. The new act replaced the confusing standard with a single, clear command.

Not surprisingly, restorative overrides are typically enacted much more quickly than updating overrides. Specifically, in our study, the average time between a restorative override and the decision itself is approximately 4.7 years, as compared to 9.5 years for a non-restorative override.¹

Depth of Override

Overrides also range in "depth"—that is, the extent to which they supersede the prior decision. Christiansen & Eskridge (2014, at p. 1533) assigned a five-point depth scale to capture this reality:

- 1 = override made a marginal change in the law;
- 2 = override overrules the point of law but leaves the prior decision in tact;

¹ Consequently, our panel data are unbalanced, a point that we address in Section 4.C.

- 3 = override overrules the point of law and the outcome in the prior decision;
- 4 = override renounces the reasoning (often indicates a restorative override);
- 5 = override renounces the reasoning and the outcome.

This variable is helpful to our study, since one would expect that citations to a case superseded by a “deep” override would decrease more than citations to a case superseded by a “shallow” override. Although depth is correlated with “type” of override, they are analytically distinct, as shown in the tabulation below (using data from our study):

Depth of Override	Restorative	Non-Restorative
1	0	15
2	0	20
3	18	79
4	19	11
5	0	2

“Conscious” Overrides

Two recent papers, Hasen (2013) and Buatti & Hasen (2015), take a different approach to identifying overrides. They include only overrides where review of the statutory language or key legislative history indicates that Congress was “consciously” responding to a prior judicial decision. Buatti & Hasen (2015) conclude that there have been 71 overrides since 1991, as compared to 122 included in Christiansen & Eskridge (2014) for the same time period. Buatti & Hasen (2015) identifies a higher percentage of restorative overrides than non-restorative overrides, suggesting (not surprisingly) that legislative history and statutory language in restorative overrides often contains specific statements referring to a disfavored judicial decision. Specifically, Buatti & Hasen (2015) includes 76% of the restorative overrides in our sample, compared to 26% of *all* overrides in our sample.

We use Christiansen & Eskridge (2014) to compile our sample because, once enacted, an override—that is, the new or newly-amended statutory law—should control resolution of a dispute whether or not Congress “knew” that a new statutory standard was superseding a prior judicial interpretation. However, we also recognize that lower courts might be more responsive to overrides where it is apparent that Congress “meant” to supersede a prior judicial decision than to overrides that could have been inadvertent or accidental. This might be particularly true if the terms of the override are at all ambiguous. Thus, we use “consciousness,” as defined by Buatti & Hasen (2015), as an explanatory variable in our regression analysis.

Overruled Cases

Our sample of overruled cases came from Brenner & Spaeth (1995) and the Supreme Court Database (SCD), originally developed by Harold Spaeth.² We searched in the database for cases that were

² The description of the category “formal alteration of precedent” in the Supreme Court Database’s codebook is consistent with—but somewhat less detailed than—the description of the factors used to identify overruled precedents in Brenner & Spaeth (1995). Compare SCD *Online Code Book*, Formal Alteration of Precedent, available at <http://scdb.wustl.edu/documentation.php?var=precedentAlteration>, with Brenner & Spaeth (1995), pp. 18-23. We reviewed the cases that were identified in the SCD for the time period covered in Brenner & Spaeth and found them to be consistent. Accordingly, we assume that the fuller description

identified as “formally altering precedent.”³ We then relied upon the research in Brenner & Spaeth (1995), and our own research of decisions for the more recent cases, to identify the precedent case(s) that the overruling cases had altered.

The category of “formally altering precedent” is defined as cases that indicate that an earlier Supreme Court decision is “overruled,” “disapproved,” “no longer good law,” “can no longer be considered controlling,” “modified and narrowed,” or that the Court “decline[s] to follow” an aspect of the earlier decision. It generally does not include cases that merely distinguish a prior precedent without indicating that the prior precedent is disapproved in some manner, at least in part (Brenner & Spaeth, 1995, p. 19-22). Brenner & Spaeth (1995) indicate that they relied primarily on majority opinions in making these assessments. However, they explain that in some instances, cases were identified as fitting within this category based on statements in dissenting opinions that a fair reading of the majority opinion was that it overruled a prior decision, even if the majority opinion itself did not admit this. Additionally, some cases were identified based on opinions in subsequent cases that indicated an earlier case had substantively overruled an even earlier decision. They note that this suggests that the universe of overruled cases in the SCD (and thus in our sample) may not be complete, in that there may be some additional cases that will later be identified as having overruled a prior precedent (Brenner & Spaeth (1995), p. 210). As stated above, Christiansen & Eskridge (2014) make the same point with respect to the overridden decisions.

We recognize that the Court engages in what has been characterized as “stealth” overruling, issuing decisions that distinguish or ignore precedents that would be expected to control the outcome of a case without acknowledging that the decision functionally overrules the prior decisions (*see, e.g.*, Eskridge (1988); Friedman (2010); Epstein et al. (2015)). The SCD presumably captures at least some of these “stealth” overrulings through its consideration of dissents and later opinions. However, many will not be included in the SCD’s “formal alteration of precedent” category and thus are not in our analysis. This could affect our study, to the extent that our sample of overrides may include what could be characterized as “stealth” overrides.⁴ We include depth and consciousness/explicitness as explanatory variables to control, at least in part, for this potential difference between our treatment groups.

The fact that some overrulings are “gradual”—in that sometimes an earlier decision narrows dramatically or undermines a prior decision and then a later decision explicitly overrules the prior

found in Brenner & Spaeth (1995) is applicable to the SCD as well, even in the years post-dating the conclusion of Brenner & Spaeth’s 1995 study, particularly since Spaeth created the database and has played a key role in maintaining it. *See SCD Primary Investigators*, available at <http://scdb.wustl.edu/about.php?s=3>.

³We chose to use the SCD/Brenner & Spaeth list of overruled cases because it has been used in several other studies (*see, e.g.*, Epstein et al. (2015); Hansford & Spriggs 2006) and because it covered our full time period. There are other sources that seek to identify all overrulings for at least portions of the time period we study, including Gerhardt (2008) and the Congressional Research Service’s *The Constitution of the United States of America, Analysis and Interpretation*. These lists have high levels of overlap with the SCD list, but they are not entirely consistent. Eskridge (1988) used a broader definition of what constitutes an “overruling” in cataloguing statutory interpretation decisions that overruled earlier statutory interpretation decisions; this study ended in 1987, near the beginning of our research window, and thus we could not use it to define our sample of cases.

⁴ Our hunch is that lower courts are relatively sensitive to such signals from the Supreme Court and likely decrease their reliance on “stealthily” overruled precedents to a considerable degree (and more than they reduce reliance on “stealthily” overridden precedents), but substantiating that phenomenon is beyond the scope of our study.

decision—has additional implications for our study. In such instances, lower courts sometimes express doubts as to which rule should apply, or circuits may split on the issue. Once the Supreme Court grants *certiorari*, lower courts may flag the issue as pending resolution. These factors may help explain why we observe in Figure 1 a gradual increase in citations to “overruled” cases in the years immediately *prior* to the event.

Depth of Overruling

As noted in the paper, overrulings, like overrides, vary in depth, ranging from cases that entirely supplant the rule announced by the earlier decision to cases that make only a minor tweak. To assess how “depth” affects subsequent citation patterns, and to permit us to better compare overrides to overruling, we developed a depth code for “overrulings”, based on the classifications used by Christiansen & Eskridge (2014) for overrides. Specifically, the depth code that we used for overrulings was as follows:

- 1 = distinguishes (or explicitly overrules) prior case on a minor point but generally follows the approach of the earlier case other than the minor disagreement;
- 2 = significantly narrows the rule of the earlier case, or identifies a broad or narrow reading of the earlier decision and chooses the narrower reading;
- 3 = distinguishes (or explicitly overrules) the reasoning of the prior case and reaches a different general rule;
- 4 = renounces the reasoning of the prior case;
- 5 = renounces the reasoning and the outcome of the prior case.

If the overruling case was a plurality decision, depth was assessed based on the concurrence that provided the necessary fifth vote, rather than the plurality opinion. We found considerable variation in “depth” for the cases in the overruled group, as shown in the table below:

Depth of Overruling	Count
1	8
2	13
3	20
4	10
5	1

Explicitness

We also sought to assess whether the explicitness of an overruling affected ongoing citation patterns. As noted above, cases were included in the SCD’s category of “formally altering precedent” based on a wide range of signals, from “overruled” to simply “declined to follow.” Additionally, although cases were identified primarily based on the majority opinions, in some instances cases were included based on dissents or statements in later opinions. An example of this is *American Trucking Ass’n v. Scheiner*, 483 U.S. 266 (1987), which is included in our dataset as overruling two earlier cases. The majority opinion suggested that it was merely distinguishing these earlier cases, whereas the dissent states that the majority opinion “directly overrules” them and argues that they should have been followed on *stare decisis* grounds (*Id.* at 298, O’Connor, J., dissenting).

To get a better understanding of how important clarity might be, we coded for whether the majority opinion, or the controlling concurring opinion if the decision did not have a majority opinion,

included an explicit statement that it was overruling a prior decision.⁵ We found that approximately two-thirds of the cases in our overruled group included such language, while one-third did not.

The explicitness of an overruling is analytically distinct from depth; in other words, some overrulings were clear that they were superseding an aspect of a prior decision but also clear that they were leaving the core holding of a prior decision in place. For example, *Quackenbush v. Allstate Insurance Co.*, 517 U.S. 706 (1996), is in our database as (partially) overruling *Thermtron Products, Inc. v. Hermansdorfer*, 423 U.S. 336 (1976). The key issue in *Thermtron* was whether a federal district court could remand a case back to state court simply because the federal court had a “crowded docket,” and a related issue of whether such a decision was subject to appellate review. The Court held a crowded docket was not a proper grounds for remand, that the district court’s decision to remand was reviewable, and that the writ of mandamus filed in the case was an appropriate vehicle for such review. In reaching this holding, the *Thermtron* Court stated in passing that an order remanding a removed action was not a final judgment that could be reviewed by appeal. *Id.* at 352. The *Quackenbush* Court explicitly “disavowed” this single statement in *Thermtron*, noting that the vehicle used for review was “peripheral” to the core issue in the case (such that it had not even been briefed) and that the statement was in tension with other precedents, *see Quackenbush*, 517 U.S. at 714-15; importantly, however, the *Quackenbush* Court also followed the more general reasoning in *Thermtron* regarding the applicable federal statutes, *see id.* at 711-12. Thus, we coded this as an explicit but shallow overruling.

An “explicit” overruling can be understood as a *rough* analog to a “conscious” legislative override. However, this is not an apples-to-apples comparison. In defining “consciousness” Buatti & Hasen (2015) looked at *both* statutory language and key legislative history. This means that lower court judges or lawyers simply looking at the statutory language of a “conscious” override would not necessarily find an “explicit” signal of disapproval comparable to an “explicit” statement that a case has been overruled.⁶ Consequently, we include both “explicit” and “conscious” as separate explanatory variables in our regression analysis, and we do not assume they will necessarily have the same effect on shadow precedent.

Overruling Cases that Were Later Overridden

There were four cases in our dataset that were identified *both* as a case that overruled a prior precedent and as a case that was itself overridden. This typically occurred when the Court changed a prior settled judicial interpretation by overruling an earlier case and Congress disagreed with the change. For example, *Griggs v. Duke Power Company*, 401 U.S. 424 (1971), established the basic proof

⁵ The vast majority of the cases we coded as “explicit” included straightforward statements that the prior decision was (at least partially) “overruled.” In some instances, we also included cases with statements that a prior precedent was “disavowed,” “disclaimed,” or “no longer controlling” where we deemed the larger discussion made it abundantly clear that this meant the prior decision was, in fact, at least partially overruled. Other signals, such as “modified and narrowed” or “rejected” might likewise be considered “explicit” evidence of an overruling, but we generally did not include them. This suggests that our “explicit” category might be somewhat underinclusive.

⁶ We separately coded for whether the name of the overridden case was included in the text of the *statute itself*. This was true for just 7% of the overrides in our sample. Even this standard is not a perfect proxy for “explicitness” of an overruling, since such statements are typically in “findings and purposes” clauses that may not be codified adjacent to the substantive language at issue in a given case.

structure used to assess “disparate impact” claims under employment discrimination statutes. In *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642 (1989), the Court substantively overruled portions of *Griggs* to make it significantly easier for businesses to defend against such claims. *Wards Cove* was a highly controversial decision, and the 1991 Civil Rights Act, Pub. L. 102-166, overrode it. *Wards Cove* is *both* an “overruling” event (with respect to *Griggs*) and an “overridden” decision (with respect to the 1991 Civil Rights Act).

Under our research design, we typically would track citations to *Griggs* for the 10 years following *Wards Cove* (from 1989-1999) to assess how quickly citations to *Griggs*, the overruled case, declined. However, the 1991 Civil Rights Act largely restored the standard initially enunciated in *Griggs*, making *Griggs*, once again, “good law.” To ensure that our results were not distorted by this subsequent Congressional action, we *only* tracked citations to *Griggs* up to 1991, the point at which Congress overrode *Wards Cove* and reinstated the *Griggs* standard. There were three other cases that fell into this category. For each, we tracked citations to the original “overruled” case only until the point at which the override of the overruling case occurred.

Appendix B: Shepard's Signal Indicators

Our analysis uses the Shepard's Signal Indicators assigned to citations by the Shepard's Citations Service, as reported by the Lexis database. Shepard's has many distinct signals that it assigns to case citations—e.g., “distinguished” or “explained” or “affirmed”. The Signal Indicators group these specific signals into broad categories as follows:⁷

Indicator	Description
Warning: Negative treatment is indicated The red <i>Shepard's</i> Signal indicator indicates that citing references in the <i>Shepard's</i> ® Citations Service contain strong negative history or treatment of your case (for example, overruled by or reversed).
Warning: Negative treatment is indicated for statute The red exclamation point <i>Shepard's</i> Signal indicator indicates that citing references in the <i>Shepard's</i> Citations Service contain <i>strong negative treatment of the Shepardized™</i> section (for example, the section may have been found to be unconstitutional or void).
Questioned: Validity questioned by citing references The orange <i>Shepard's</i> Signal indicator indicates that the citing references in the <i>Shepard's</i> Citations Service contain treatment that questions the continuing validity or precedential value of your case because of intervening circumstances, including judicial or legislative overruling.
Caution: Possible negative treatment indicated The yellow <i>Shepard's</i> Signal indicator indicates that citing references in the <i>Shepard's</i> Citations Service contain history or treatment that may have a significant negative impact on your case (for example, limited or criticized by).
Positive treatment indicated The green <i>Shepard's</i> Signal indicator indicates that citing references in the <i>Shepard's</i> Citations Service contain history or treatment that has a positive impact on your case (for example, affirmed or followed by).
Citing references with analysis available The blue “A” <i>Shepard's</i> Signal indicator indicates that citing references in the <i>Shepard's</i> Citations Service contain treatment of your case that is neither positive nor negative (for example, explained).
Citation information available The blue “I” <i>Shepard's</i> Signal indicator indicates that citing references are available in the <i>Shepard's</i> Citations Service for your case, but the references do not have history or treatment analysis (for example, the references are law review citations).

Home > All About... > Using the *Shepard's*® Citations Service > The *Shepard's* Signal™ Indicators

We did not independently assess the reliability of the signal indicators, but Spriggs & Hansford (2000) evaluated their accuracy and found them to be generally reliable, with the stronger negative treatment codes being the most reliable. Accordingly, we believe that the distinction between “positive” cites and “negative” cites that we use in our analysis likely captures real differences in how courts cite to prior precedents.

⁷ This chart may be found at: http://help.lexisnexis.com/tabula-rasa/newlexis/shepardsignal_ref-reference?lbu=US&locale=en_US&audience=all,res,cb,lps,med,vsa,tax,lpa,icw,blink,pub,urlapi.

That said, Spriggs & Hansford (2000) only assessed reliability of citations in subsequent Supreme Court decisions; it may be that Shepard's has higher levels of quality control for accuracy in coding Supreme Court decisions than lower court decisions. Additionally, the interpretive complexity implicit in integrating overrides into analysis of precedent may result in a higher level of false positives and false negatives than is typical. (As discussed in Widiss (2014), Christiansen and Eskridge (2014) study of overrides, which uses Westlaw flags, uncovered relatively high numbers of false positive and false negative flags.) Despite these limitations, use of Shepard's signals was necessary to assemble a significantly large body of data.

Appendix C: Headnote Coding Protocol

As explained in the text, we used LexisNexis “headnotes”, and the possibility of Shepardizing by headnote, to gain a clearer understanding of *which* propositions within a given case were being relied upon by later cases. 60 cases that were *overridden* and 20 cases that were *overruled* were randomly selected for this analysis. This randomized list was created using Excel’s random number generator—function: =RAND().

After identifying cases for analysis, we read the case, the override statute or overruling case, and, in some cases, additional later cases or explanatory materials to assess how the override statute or overruling case interacted with the precedent case. Then, for each of the randomly selected cases, each LexisNexis “headnote” was initially coded using 5 numerical classifications, although, as described below, we ultimately combined two categories for analysis and discarded one category for analysis, yielding the 3 numerical classifications that we address in the text:

1—Propositions that were directly superseded by the new statute or directly overruled by the later decision. Headnotes were coded as a “1” if *any aspect* of the statement in the headnote would no longer be correct under the new governing standard established by the override statute or the overruling case. Sometimes a single headnote would include both statements that were superseded and statements that were not addressed by the override or overruling.

2—Reasoning supporting the proposition that was overridden or overruled but not squarely superseded by the text of the override or overruling opinion. This category now includes assertions of general interpretive methodologies (e.g., canons of statutory interpretation) that were cited in support of the proposition that was subsequently overridden. Originally, such canons were separately coded.

3—Issues addressed in the case that were totally irrelevant to the override or overruling (e.g., procedural issues, other statutory provisions, etc.). This included interpretive methodologies that supported propositions irrelevant to the override or overruling, as well as citations to statutory provisions that were totally unrelated to the override or overruling.

When we originally coded the headnotes, we included a fourth category that was used for very general descriptions of the statute at issue in the override or overruling; paraphrases or quotations of relevant statutory language; and headnotes that simply listed the statutory provision interpreted in the proposition that was superseded without additional text. Because many of these headnotes were simply to statutory provisions without discussion, and often the statutory provision remained the same even after an override had been enacted, or the overruling case had been decided, it was unclear whether subsequent citations to these headnotes were functionally citing the prior precedent or the new rule. Accordingly, we dropped them from our analysis.

Two aspects of this coding protocol require further elaboration. With some regularity, later statutory amendments added language to a statute that was within plausible meaning of the pre-existing language. For example, *Bailey v. United States*, 516 U.S. 137 (1996), concerned whether 18 U.S.C. § 924(c)(1), which provided mandatory minimum sentences for “use” of a firearm in relation to certain crimes, encompassed possession of a firearm. Circuits had split on the issue. The Court interpreted “use” to require more than mere possession. Congress subsequently amended the relevant provision to explicitly permit the enhancement for “possess[ion] of a firearm” “in furtherance of any such crime.” The headnote in *Bailey* stating that “use” must connote more than mere possession was coded as a “1” since possession would now lead to the sentencing enhancement, even though the Court’s interpretation of “use” as requiring more than “possession” arguably remains controlling.

Additionally, in some instances, the statutory provision actually interpreted in the overridden case remained unchanged, but Congress passed a new statute or added a new provision that had the effect

of changing the result at issue in the case. For example, *United States v. Kozminski*, 487 U.S. 931 (1988), concerned the meaning 18 U.S.C. § 1584, which prohibits knowingly and willfully holding an individual in “involuntary servitude.” The Court interpreted this phrase to require a showing of use or threat of physical harm or legal coercion, rather than use of other forms of coercion such as psychological coercion. The Court drew this meaning from prior interpretations of the 13th Amendment, which also uses the phrase “involuntary servitude.” The override added a new provision, 18 U.S.C. § 1589, that criminalizes forced labor through physical, legal, or other psychological, emotional, or reputational coercion, as a means of criminalizing the kind of behavior that had been at stake in *Kozminski*. Although 18 U.S.C. § 1584’s meaning may not have technically changed, the new provision encompasses the conduct at issue in the case being interpreted, and the headnotes discussion 18 U.S.C. § 1584 as requiring a showing of physical or legal coercion rather than other forms of coercion were coded as “1s”. Headnotes discussing the 13th Amendment’s meaning, by contrast, were coded as “2s”, since the override cannot change the meaning of the constitutional term.

This research design was chosen because it let us efficiently sort later citations according to which principles in the overridden or overruled case were relied upon. That said, it was obviously dependent on—and therefore in some ways limited by—LexisNexis’s editorial choices regarding headnotes. We chose to use LexisNexis’s headnotes rather than Westlaw’s key cites because a review of a sample of cases suggested that LexisNexis’s headnotes were usually more detailed than Westlaw’s, and LexisNexis seems more frequently to include interpretative methodologies in the headnotes. That said, there was wide variability in the detail in which LexisNexis assigned headnotes to the cases in our dataset. For example, some cases had as few as one headnote, while others had as many as twenty-eight, with the variation only partially explained by differences in the numbers of issues addressed in cases.

In a small number of cases in our dataset, none of the headnotes in the case articulated the proposition that was later superseded. Thus, in some cases in our dataset, none of the headnotes were assigned a “1”, though there were always “2”s that were related to the superseded proposition. In many cases, LexisNexis editors identified in headnotes numerous different interpretative methodologies as well as substantive application of those principles. But in some cases, including other cases which likewise used several interpretative methodologies, none of the interpretative methodologies were included among the headnotes. This was true even for some cases, such as *Reno v. Bossier Parish*, 528 U.S. 320 (2000), which are frequently cited for the interpretative methodologies employed. The opposite was true as well; in a few instances, e.g., *United States v. Nordic Village Inc.*, 503 U.S. 30 (1991), the headnotes all focused on interpretative methodologies and none delved into the substance of the actual statute at stake in the case. Thus, our analysis derived from shepardizing by headnotes cannot capture entirely the various principles for which the cases in our dataset are cited.

In general the scale worked relatively well, but coding sometimes required difficult judgment calls regarding precisely what was superseded (affecting the lines between “1” and “2”), as well as where the lines between “specific supporting reasoning” (“2”) and general background (the category we later discarded) should be drawn. Occasionally, there was also some level of judgment call required to distinguish “irrelevant” (“3”) from “background reasoning” (“2”). We found that it particularly difficult to code cases concerning *habeas corpus* on this scale, because *habeas* cases typically relied heavily on a detailed and extensive background of precedent that concerned largely constitutional principles and prior decisions.

Fourteen overridden cases were coded by two individuals to determine a measure of inter-coder reliability. One case was used as model that we walked through together. We then each independently coded an additional nine, met to discuss these cases, and then independently coded an additional four, which we also met to discuss.

There were 171 headnotes coded by both individuals. Using our revised scale that combined interpretative methodologies with other supportive reasoning, but includes the fourth category that we ultimately discarded from analysis, there was complete agreement for 73.7% of the footnotes (using the five-point scale we employed originally, there was complete agreement on 71.3% of the headnotes). There was disagreement by 1 number (e.g., one would code the proposition a “1” and one would code the proposition a “2”) for 24% of the headnotes. There were only three headnotes (1.7%) on which the coding differed by two numbers. In each instance in which we disagreed as to the classification, we met, discussed the issue, and then jointly decided which classification we thought was appropriate. In many instances, the headnotes on which we disagreed were headnotes that one or both of us had flagged as “on-the-line” between two numbers. In some instances as well, the “same” issue would explain differences on multiple distinct headnotes within a given case. For example, *Wilson v. Garcia*, 471 U.S. 261 (1985), concerned interaction of 42 U.S.C. § 1983 and state law, and three headnotes simply referenced the particular New Mexico statute at issue in the case. One of us coded all three as irrelevant to the interpretation of the federal statute, while the other one of us coded all three as general background (the category that we subsequently excluded from analysis). This was counted as three different headnotes on which our coding differed.

Appendix D: STATA commands for replication of regression results

To allow replication of our regression results, we have posted, as supplemental materials on the Journal of Legal Studies website, two data files:

- Case Level Data - Sept 2016.dta
- Headnote Level Data - Sept 2016.dta

The first file – Case Level Data - Sept 2016.dta – can be used to estimate the regressions models in Tables 4 and 7. The second file – Headnote Level Data - Sept 2016.dta – can be used to estimate the regression models reported in Table 6.

This appendix provides STATA commands for replicating each model reported in the paper. Except as indicated in the notes below, variable definitions are included in the text of the paper.

Regression commands for Table 4⁸

```
use "Case Level Data - Sept 2016.dta"
```

M1

```
xtreg NLogCitations post post_override post_overruled Log_YSD i.citingyear  
if deadzone12==0 & yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 & t!=0  
& PrePostData==1, fe vce(robust)  
test post_override = post_overruled
```

M2

```
xtreg NLogCitations post post_override post_overruled depth_x_post_all  
Restorative_p negativeDV_Learn123_post positiveDV_Learn123_post  
sc_post_net citations sc_post_warning Log_YSD i.citingyear if deadzone12==0  
& yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 & t!=0 &  
PrePostData==1, fe vce(robust)  
test post_override = post_overruled  
test post_override + Restorative_p = post_overruled
```

M3

```
xtreg NLogCitations post post_override post_overruled depth_x_post_all  
Restorative_p conscious_p explicit_p negativeDV_Learn123_post  
positiveDV_Learn123_post sc_post_net citations sc_post_warning Log_YSD  
i.citingyear if deadzone12==0 & yearsbetween!=0 & mean_pre_z > 3 &  
first_year!=1 & t!=0 & PrePostData==1, fe vce(robust)  
test post_override = post_overruled  
test post_override + Restorative_p = post_overruled  
test post_override + conscious_p = post_overruled + explicit_p
```

⁸ The variable “deadzone12” equals 1 if t=1 or t=2, otherwise deadzone12 = 0. The variable “yearsbetween” equals the number of years between the original decision and the superseding event. The variable “PrePostData” is a dummy variable which equals 1 if the case includes observations at t=-1 and t=3 and otherwise is set to zero. The variable “mean_pre_z” equals the average number of citations that the case received per year in the pre-event period. The panelvar for purposes of STATA’s xtset command is “spaeth_panel”.

M4

```
xtreg NLogCitations post post_override post_overruled Log_YSD i.citingyear
if deadzone12==0 & yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 &
matched==1 & t!=0 & PrePostData==1, fe vce(robust)
test post_override = post_overruled
```

M5

```
xtreg NLogCitations post post_override post_overruled depth_x_post_all
Restorative_p negativeDV_Learn123_post positiveDV_Learn123_post
sc_post_net citations sc_post_warning Log_YSD i.citingyear if deadzone12==0
& yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 & matched==1 & t!=0 &
PrePostData==1, fe vce(robust)
test post_override = post_overruled
test post_override + Restorative_p = post_overruled
```

M6

```
xtreg NLogCitations post post_override post_overruled depth_x_post_all
Restorative_p conscious_p explicit_p negativeDV_Learn123_post
positiveDV_Learn123_post sc_post_net citations sc_post_warning Log_YSD
i.citingyear if deadzone12==0 & yearsbetween!=0 & mean_pre_z > 3 &
first_year!=1 & matched==1 & t!=0 & PrePostData==1, fe vce(robust)
test post_override = post_overruled
test post_override + Restorative_p = post_overruled
test post_override + conscious_p = post_overruled + explicit_p
```

Regression commands for Table 6⁹

```
use "Headnote Level Data - Sept 2016.dta"
```

M7

```
xtreg NLogCitations post post_c1 post_c2 i.citing_year if deadzone12==0 &
yearsbetween!=0 & first_year==0 & mean_pre_z > 2 & hn_c3==0 & t!=0, fe
vce(robust)
test post_c1 = post_c2
```

M8

```
xtreg NLogCitations post post_c1 post_c2 Overruled_post Overruled_C1_post
Overruled_C2_post i.citing_year if deadzone12==0 & yearsbetween!=0 &
first_year==0 & mean_pre_z > 2 & hn_c3==0 & t!=0, fe vce(robust)
test post_c1 = post_c2
test Overruled_post + Overruled_C1_post = 0
```

⁹ The variable "hn_c3" is a dummy variable that equals 1 if the headnote provides a general descriptions of the statute at issue in the override or overruling or if the headnote only listed the statutory provision interpreted in the proposition that was superseded without additional text. As explained in Appendix C, these headnotes are excluded from the analysis. For the headnote analysis, the panelvar for purposes of STATA's xtset command is "usc_HN_partial"

Regression commands for Table 7

```
use "Case Level Data - Sept 2016.dta"
```

M9

```
xtreg NewDV4 post post_override post_ouerruled depth_x_post_all  
Restorative_p negativeDV_Learn123_post positiveDV_Learn123_post  
sc_post_net citations sc_post_warning Log_YSD i.citingyear if deadzone12==0  
& yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 & t!=0 &  
PrePostData==1, fe vce(robust)  
test post_override = post_ouerruled
```

M10

```
xtreg NLogCitations post post_override post_ouerruled depth_x_post_all  
Restorative_p negativeDV_Learn123_post positiveDV_Learn123_post  
sc_post_net citations sc_post_warning Log_YSD i.citingyear if deadzone12==0  
& t!=3 & t!=4 & t!=5 & yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 &  
t!=0 & PrePostData==1, fe vce(robust)  
test post_override = post_ouerruled
```

M11

```
xtreg NLogCitations post post_override post_ouerruled depth_x_post_all  
Restorative_p negativeDV_Learn123_post positiveDV_Learn123_post  
sc_post_net citations sc_post_warning Log_YSD i.citingyear if t!=0 &  
yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 & t!=0 & PrePostData==1,  
fe vce(robust)  
test post_override = post_ouerruled
```

M12

```
xtreg NLogCitations post post_override post_ouerruled depth_x_post_all  
Restorative_p negativeDV_Learn123_post positiveDV_Learn123_post  
sc_post_net citations sc_post_warning Log_YSD if deadzone12==0 &  
yearsbetween!=0 & mean_pre_z > 3 & first_year!=1 & t < 4 & t > -2 & t!=0 &  
PrePostData==1, fe vce(robust)  
test post_override = post_ouerruled
```

M13

```
xtreg NLogCitations post post_override post_ouerruled depth_x_post_all  
Restorative_p negativeDV_Learn123_post positiveDV_Learn123_post  
sc_post_net citations sc_post_warning unified_control Fed_id_aligned_Over_p  
liberal_case_post Log_YSD i.citingyear if deadzone12==0 & yearsbetween!=0  
& mean_pre_z > 3 & first_year!=1 & t!=0 & PrePostData==1, fe vce(robust)  
test post_override = post_ouerruled
```

Appendix References

- Brenner, Saul, and Harold J. Spaeth. 1995. *Stare Indecisus: The Alteration of Precedent on the Supreme Court, 1946–1992*. New York: Cambridge University Press.
- Buatti, James, and Richard L. Hasen. 2015. Conscious Congressional Overriding of the Supreme Court, Gridlock, and Partisan Politics. *Texas Law Review See Also* 93:263–88.
- Christiansen, Matthew R., and William N. Eskridge, Jr. 2014. Congressional Overrides of Supreme Court Statutory Interpretation Decisions, 1967–2011. *Texas Law Review* 92:1317–1515.
- Epstein, Lee, William M. Landes, and Adam Liptak. 2015. The Decision to Depart (or Not) from Constitutional Precedent: An Empirical Study of the Roberts Court. *New York University Law Review* 90:1115–59.
- Eskridge, William N., Jr. 1988. Overruling Statutory Precedents. *Georgetown Law Journal* 76:1361–1439.
- . 1991. Overriding Supreme Court Statutory Interpretation Decisions. *Yale Law Journal* 101:331–455.
- Friedman, Barry. 2010. The Wages of Stealth Overruling (with Particular Attention to *Miranda v. Arizona*). *Georgetown Law Journal* 99:1–63.
- Gerhardt, Michael J. 2008. *The Power of Precedent*. New York: Oxford University Press.
- Hansford, Thomas G., and James F. Spriggs II. 2006. *The Politics of Precedent on the U.S. Supreme Court*. Princeton, NJ: Princeton University Press.
- Hasen, Richard L. 2013. End of the Dialogue? Political Polarization, the Supreme Court, and Congress. *Southern California Law Review* 86:205–61.
- Spriggs, James F., II, and Thomas G. Hansford. 2000. Measuring Legal Change: The Reliability and Validity of *Shepard's Citations*. *Political Research Quarterly* 53:327–41.
- Widiss, Deborah A. 2014. Identifying Congressional Overrides Should Not Be So Hard. *Texas Law Review See Also* 92:145–69.