# Courting the Academy
## The Judicial Role in Popularizing Legal Scholarship*

Matthew Estes† Ransi Clark‡

May 25, 2025

**Abstract**

The relationship between the courts and legal academia is widely studied. Although there is much work on how the courts *use* academic legal theories, how courts *affect* the production, development, reception, and canonization of academic legal theory is under-explored. This paper studies how the courts impact the popularity of certain legal theories. Combining data on academic citations and judicial mentions, this paper quantifies how judges impact the popularity of two legal movements: legal realism and originalism. We find that judicial mentions increased academic citations to legal realist works, a finding that was robust to model specifications. For originalists, the sign of the effect (increase/decrease) is most often not identified and the magnitude is smaller.

This paper also contributes to research methodology for empirical legal studies. In settings with long treatment histories, counterfactual identification is complicated when unit-level treatment effects depend on treatment timing or even the entire treatment history. We propose an identification strategy to bound (partially identify) causal effects in such a setting. The partial identification strategy we propose uses information from when units experienced the same treatment (pre-divergence) to bound counterfactuals for those units when their treatment paths differ (post-divergence). Throughout the methods section we provide examples and illustrations to explain causal inference problems and our bounding strategy.

**Keywords:** empirical legal studies, courts, realism, originalism, counterfactuals
**Word Count:** 16,149

*This research uses data made available through JSTOR's Constellate platform, operated by ITHAKA. The findings and conclusions are solely those of the authors. We are grateful to Mike Alvarez, Jonathan N. Katz, Ryan T. Moore, Ronak Patel, Bob Sherman, and participants at the 2025 MPSA Conference for helpful comments and feedback on this project.

†California Institute of Technology. mestes@caltech.edu

‡California Institute of Technology. ransi@caltech.edu

# 1 Introduction

The Supreme Court's recent opinion in *Franklin v. New York*[1] exemplifies the present relationship between the judiciary and the legal academy. In *Franklin*, the Court denied *certiorari* to reconsider its Confrontation Clause jurisprudence under the Sixth Amendment. Yet Justices Alito and Gorsuch issued opinions signaling a willingness to overturn the leading Confrontation Clause case.[2] Why? Because, Justice Alito writes, "recent scholarship ... casts doubt on key aspects of *Crawford's* reasoning."[3]

The Court's encouragement of particular legal scholarship (esp. originalist scholarship) is not new. Indeed, a not insubstantial amount of empirical scholarship studies how and when judges use (or ignore) traditional legal scholarship. But much less study is devoted to the reverse relationship of how and when the legal academy responds to judicial incentives. Indeed, despite criticisms that the legal academy is disengaged with the practical matters of lawyering and the judiciary, individual academics hold judicial mentions of their work in high esteem. Academics with judicial ambitions are incentivized to produce work that they believe the judiciary will invoke. To provide for the judiciary, academics must know what the judiciary needs. Judges can signal a need for further academic work by mentioning certain legal pieces in their opinions. Other judges may produce the scholarship themselves.[4] Still others may convey signals by indirect means, e.g. through their clerks who later become academics.

The purpose of this paper is quantification of the extent to which the academy heeds the judiciary's call. More specifically, we consider the courts' role in the rise of realist scholarship in the early 20th century and the rise of originalist scholarship in the late 20th century. We begin by identifying scholarly works (books, chapters, or articles) that are most relevant to legal realism or originalism. Then we obtain the judicial mentions[5] and academic citations for each work. Those specific series are obtained with a view to estimating how much of the scholarly attention these works get is attributable to court mentions. Although citations are an imperfect measure of scholarly attention, they help us to not just measure popularity in academic circles but also the volume of academic production.

A causal relationship between judicial mentions and related scholarly production

[1] *Franklin v. New York*, 604 U.S. (2025).

[2] *Crawford v. Washington*, 541 U.S. 36 (2004).

[3] *Franklin v. New York*, 604 U.S. (2025) (Alito, J., concurring).

[4] Consider Judges Posner and Easterbrook, or Justice Scalia.

[5] We use the term "judicial mentions" instead of the more familiar "judicial citations" to avoid confusion between citations arising from the judiciary and the academy.

is difficult to pin down because the two series may be simultaneously affected by the politics of judicial appointment.[6] Academics that see the political tides changing may get ahead of it by producing scholarly work that is adherent to a particular judicial movement.[7] Despite these difficulties, there is much variation in how works are mentioned by the judiciary. Some works are mentioned much more often than others.[8] Variation also exists in how the time since publication before a work is mentioned by the judiciary. These variations aid us in causal analysis by providing stand-ins for the realized citation series under an alternative scenario.

These challenges to causal arguments are encountered across the social sciences. Because of the insolubility of many such challenges, there is a preference towards designing and implementing experiments whenever possible. Unfortunately, the question at hand is not easily amenable to experimentation. Such is the case for most questions that concern the judiciary, not least because diversions from normal operations even for scientific purposes may be seen as a miscarriage of justice. Yet, the impossibility of experimentation should not be fatal to the study of these processes, provided the counterfactual assumptions underpinning the interpretation of the derived estimates as causal hold.

The methodology we use to impute counterfactuals is simple and easily interpreted. The method, which we call the bounded deviations approach, allows for a wide range of counterfactual outcomes for the relative difference between the citation trajectories. The intuition is that, absent treatment, treated units would've experienced an outcome "close to" or "nearby" the control unit outcomes.

Counterfactual methods require us to identify (albeit somewhat imperfectly) a set of treatment (works that were mentioned) and control (works not mentioned).[9] The citation outcome of the control set is used to construct counterfactual citations for the treatment set under alternative mention paths.[10] What we consider treatment

[6]Both realist and originalist movements were colored by the political leanings of the Administrations that nominated adherent judges to the court. The Roosevelt Administration nominated several legal realists and the Reagan Administration nominated originalists.

[7]They may also do so in anticipation of a judicial appointment or a blockbuster court case.

[8]Judges and their clerks may also use, for instance, citation counts as a heuristic in determining whether the work should be mentioned in their opinion or not. This potentially summons the pernicious problem of reverse causality. The possibility is not that the judicial mentions increased the scholarly popularity of the work, but its own baseline level of scholarly popularity.

[9]The process requires coarsening of the paths since some works receive a disproportionately large number of mentions. There may also be too few works that receive no mentions from the judiciary, even though there will be a large number of works that receive a very small number of mentions dispersed over several decades.

[10]The textbook version of the differences-in-differences strategy encodes treatment as a singular

here, the mentions by a court, is an entire sequence of events, the impact of which is not immediate within the year. A treatment that is a sequence of events rather than one event can cascade over time too. We accommodate these idiosyncrasies by taking a flexible approach to ensure that the best comparisons are made.[11]

We find that judicial mentions increase academic citations to realist works. These effects generally appear with some delay and are relatively robust to the identification tuning parameter we use. We show, for example, that realist works first mentioned in judicial opinions from 1976–2025 received a positive boost in academic citations. Although the magnitude of the effect is not identified, it is between 5–25 added academic citations across model specifications.

The impact of judicial mentions on originalism is less certain. Because it is a nascent judicial theory, it is too early to tell whether academic citations to originalist works increase or decrease in response to judicial mentions. The lag between issued judicial opinions and the development or notice of academic legal theories suggest the data will not pin down the effect for some time. Evidence of causality (or lack thereof) may, in other words, be premature because an uncertain amount of time may need to elapse for an academic work to trickle into the judicial milieu.[12]

Our work contributes to several strands of literature. The most established of these is that of the origin of jurisprudential movements itself. We bring citation metrics to bear on the evolution of such movements. The beginnings of jurisprudential movements, in particular those that have partisan bearings, are usually prone to revision[13] and reinvention.[14] The process of identifying influential texts and then tracing their influence on the academy and the judiciary is less prone to such temptations.

Another strand of literature, beginning with Judge Harry Edwards' observation that "law and" movements are displacing practical legal training in law schools,[15] seeks to quantify the relationship between the legal academy and the judiciary with a view to studying the divergence between the two institutions. The vast majority

event—such as the raising of the minimum wage, or the enactment of a gun law—whose impact is felt immediately.

[11]The approach is flexible compared to the regression implementation for the textbook differences-in-differences. Yet, limitations exist even within this more flexible framework because the more varied the treatment paths are the more numerous their possible counterfactuals become. Statisticians call this problem the curse of dimensionality.

[12]Because little of the mechanism by which judges learn of legal scholarship is known (except for their own pronouncements), it is difficult to know when it is too early to assess the causal impact of judicial (non-)mention on the fate of the publication.

[13]See Tamanaha (2008) for realism and Sawyer III (2018) for originalism.

[14]Both legal realism and originalism has "new"-er versions.

[15]Edwards (1992)

of such empirical studies merely quantify the volume of judicial mentions of legal scholarship. While volume is an important metric and adequate for an exploration of whether the judiciary uses legal scholarly work, it does little in the way of illuminating whether academics respond to the judiciary. By focusing on particular substantive movements, we are better able to measure the effect of judicial mentions on subsequent scholarly work relevant to that movement.

Lastly, this study exemplifies the utility of bounding (partial identification) in applications where there is great uncertainty as to whether the assumptions required to make causal claims are satisfied. This approach—most widely associated with econometrician Charles Manski[16]—relaxes counterfactual assumptions to bound the possible causal effects within informative ranges. In this vein, the bounded deviations method we use is robust to violations of assumptions required for causality with other approaches. For example, our approach nests the parallel trends assumption for the difference-in-differences ("DID") strategy.[17] In the DID approach, comparisons are made between citation trajectories re-centered at their baselines. This can inoculate against a variety of challenges to causal claim-making such as: popularity of the academic work (or the author) before judicial mentions, citation trends that are affected by factors extraneous to judicial mentions, or citation trends related to the time elapsed since publication.

The remainder of the paper is structured as follows. In Section 2, we review existing empirical and non-empirical work on how courts use academic legal theories and works. We also discuss how existing scholarship largely neglects how the courts affect legal academics. In Section 3, we give an overview of two important U.S. legal movements: American Legal Realism and Originalism. We discuss the observational data used in our study and some high-level descriptive findings from the data. In Section 4, we discuss the causal inference problem in this setting. We describe simplifying assumptions and illustrate how our causal identification strategy works. Finally, in Section 5, we apply our methodology to estimate the impact of judicial mentions on academic citations to realist and originalist works.

## 2 Legal Theories in the Academy and in the Courts

The relationship between the legal academy and the courts is a subject of interest for many legal scholars. Still, there is much descriptive empirical work that remains

[16]See, e.g., Manski (1997), Manski (2007), and Manski and Pepper (2018)

[17]See also Manski and Pepper (2018).

undone. Although there has been much hand-wringing about how academics affect (or not) judges and the judicial process, the reverse channel is relatively underexplored. In particular, credible empirical evidence of the causal effects of judicial activities on the development, production, and reception of academic legal theories is not well-known. This section discusses prior work—both empirical and non-empirical (qualitative, anecdotal)—on how the courts and legal academy interact with one another.

### 2.1 How Courts Use Academic Legal Articles

One well-known lament is that legal academics are largely irrelevant to judges and judging. There is a smattering of qualitative and anecdotal evidence for this claim. Consider, for example, Judge Posner's claims along these lines:

> "So there is much to criticize in the judicial profession and therefore much room for improvement. But where is the improvement to come from? A possibility that appeals to me as a former law professor is the law schools. Law professors write a great deal about the judiciary—and mainly the federal judiciary. But there is a question about how well informed about, or helpful to, the judiciary that writing is. At present, not very, I have discovered. Not that I'm inclined to apply Brendan Behan's comment about film critics to law professors by comparing the professors (relative to judges) to 'eunachs in a harem; they know how it's done, they've seen it done every day, but they are unable to do it themselves.' Some judges might think the comparison apt, however."[18]

Empirical findings support this hypothesis. At the Supreme Court, for example, Newton (2012) finds that only 20.19% of signed opinions mention a law review article from 2001-2011.[19] This may, in part, stem from ideological differences between the justices and the authors of law review pieces: the three justices with the highest percentage of opinions mentioning law reviews were liberal justices (Breyer 26.14%, Stevens 24.54%, and Ginsburg 23.08%), whereas the three justices with the fewest citations to law reviews were all conservative justices (Thomas 13.31%, Rehnquist 13.79%, and Roberts 16.04%).[20]

[18] Posner (2016).
[19] See Table 1 in Newton (2012).
[20] See Table 1 in Newton (2012).

This trend may be changing. Although the number of mentions to law review pieces is generally low,[21] there is some evidence that law review pieces are increasingly mentioned, especially elite law reviews in high-profile Supreme Court cases.[22] Indeed, Detweiler (2020) finds that, although mentions of law review articles remained low in 2018 (1.8%), opinions increasingly mention law reviews relative to the historical minimum in 2009 (1.47%).[23] Note that the Detweiler (2020) study uses the entire universe of reported opinions from U.S. state and federal courts 1945–2018.

What is driving this possible shift remains unclear. One possibility is that law review articles are becoming more sophisticated and empirical, which can "bolster a doctrinal claim".[24] Indeed, the amount of "minimal empirical content" in law review articles increased from 1998-2008. [25]

Another important factor affecting judicial mentions of academic (and non-academic) sources is technological progress. Fronk (2010) analyzes federal appellate citation patterns and finds that the average number of cited cases increased from 15.66 in 1957 to 31.14 in 2007.[26] This was a period of technological progress, which substantially lowered the cost and difficulty of locating relevant legal sources. A similar effect may be relevant for citations to law review works, which are increasingly easy-to-access and widely shared, if law review pieces are *complements* to primary legal sources.[27]

A final consideration is that legal opinions may cite law review articles because they speak to a judge's own theories. Legal theories may not constrain judicial decisions writ-large, but they may persuade an individual judge to the extent they align with the judge's own theories.[28] Indeed, judicial mentions of more jurisprudential or theoretical may represent alignment with or acceptance of theoretical precepts, means, and ends. To the extent law review works better reflect the theories of practicing judges, we may expect more citations to corroborating law review works.[29]

[21] Sirico Jr. and Drew (1991)

[22] See Feldman (2018).

[23] Detweiler (2020) attributes the overall decline from the peak in early 1980s to substitution with primary legal sources.

[24] Diamond and Mueller (2010) at 595.

[25] See Figure 3 in Diamond and Mueller (2010).

[26] See Fronk (2010) at Table 1.

[27] But see Detweiler (2020) for discussion of law review works and primary legal sources being *substitutes*. In this case, the technological progress for easily accessing primary legal materials would decrease citations to law reviews. Indeed, Detweiler (2020) argues that: "By facilitating the discovery of primary law, these services provided researchers with a viable, and in many ways superior alternative to using law reviews as a shortcut for their own research."

[28] See, e.g., D'Amato (1999); see also D'Amato (1989).

[29] If this explanation is correct, should we expect further increases in citations to law reviews as machine learning and generative language models perform increasingly well at legal reasoning tasks

## 2.2 How Courts Affect Legal Academia

Even if judges find law reviews or legal academia "irrelevant", this does not imply that judges are irrelevant to law professors. Indeed, judges may impact the development in many ways, intentional or otherwise.

The most obvious influence of the courts on legal academia is on legal educational materials. Because of the casebook method, judges impact legal education through written opinions. Empirical descriptive work suggests that certain superstar judges (e.g. Judges Richard Posner and Frank Easterbrook) dominate the casebooks.[30] Longer service (i.e. seniority) is also positively correlated with casebook inclusion.[31] Other studies suggest that additional factors may explain casebook inclusion, especially for non-senior judges with a high degree of casebook inclusion.[32]

Another way judges can impact law professors is by changing the demand for certain legal theories. Simple economic incentive arguments imply that academics respond to the non-pecuniary benefits—e.g. prestige, fame, citations, interest—from mentions in court opinions. Indeed, increased court demand (support) for particular legal theories should, on balance, increase the supply of those theories unless law professors and judges operate in fundamentally different "markets.". Given the influx of scholars from other disciplines into the legal academy, academics may feel more pressure to impress like-minded colleagues rather than communicate to judges on the practical matters of the law. For some judges and some law professors, this is certainly a possibility, as Judge Posner notes:

> "[T]here is a wild literature that I have avoided mentioning in which law professors in immensely long articles subject legal texts to the hermeneutic techniques of postmodernist literary theory. No judge could get *anything* out of that literature, and this unbridgeable gap is not merely a generational one."[33]

Judge Posner offers an especially vivid potential explanation for this phenomenon:

> "The process is Darwinian. In nature each animal species must find a niche for itself, critically including a food that it can find and eat with-

and law professors (and lawyers) are better able to pin down a judge's operant theories? See, e.g., Thalken et al. (2023) (discussing LLMs and legal reasoning tasks).

[30] Gulati and Sanchez (2002).

[31] Gulati and Sanchez (2002).

[32] See, e.g., Fitzpatrick and Varghese (2017).

[33] Posner (1992) at 1928.

> out encountering destructive competition from another species. In the academy each species of professor must find an academic niche in which he can avoid destructive competition from other professors. ... Their need to communicate with persons outside of their niche, such as judges, like the need of a squirrel to learn to eat dandelions as well as nuts, is minimized."[34]

In other words, there is nothing inherently "wrong" with market specialization: judges and law professors simply occupy different niches.

Politics and ideology are also likely relevant to whether legal scholars respond to courts. Because legal scholarship is often ideological,[35] conservative legal scholars may write, for example, to respond to the current Supreme Court whereas liberal legal scholars will not. In other words, scholars may operate in separate ideological silos, which can affect whether certain law professors respond (or not) to court demand.

The focus of the remainder of this paper is testing whether law professors do, in fact, respond to judicial mentions. We examine whether certain academic law pieces increase in popularity following mentions in published court opinions.

## 3 Data Collection & Descriptive Findings

This paper uses data on two important legal movements: American Legal Realism and Originalism. We study these two legal movements because of their historical and contemporary impact on the law. As Posner (2016) notes, realism and originalism are two of the most influential legal philosophies in the United States:[36]

> "Legal realism...was a highly influential judicial philosophy not just in the 1920s and 1930s but in the entire period that began with the publication of Holmes's *The Common Law* in 1881 and ended with the end of the Carter Presidency a century later. Reagan's election, his conservative appointments to the federal courts of appeals and the Supreme Court, and the more or less simultaneous creation of the Federalist Society, began to shift the federal judicial balance back toward formalism, giving rise to 'originalism' and 'textualism' and increased resort to dictionaries ('liter-

[34] Posner (2016) at 8-9.
[35] See, e.g., Table 3 and Figure 3 in Chilton and Posner (2015).
[36] As Posner (2016) notes, the origins of these philosophies were also political.

alism') as sources of statutory meaning—all backward-looking sources of judicial guidance."[37]

Because of the ideological character of these movements, their endorsement was likely to bring greater rewards.[38] Judges are likely to signal their endorsement by mentioning prominent scholarship of the movement. The mention of these by the courts can have an encouraging effect on legal academics who seek to join the fray.

To assess the impact of courts on these two movements, we collected two main types of data for each movement. The first type of data is data on judicial mentions, that is, how often judges cite specific legal works (articles) for thinkers associated with the movement. The second type of data is information on the popularity of different legal works, as measured by academic citations. The data sources, collection process, and high-level descriptive findings are summarized for each movement below.

## 3.1 Legal Realists

### 3.1.1 Description of the Movement

American Legal Realism greatly impacted the development of the law in the United States. For legal education and scholarship, the realists' primary contribution was in developing a non-formalist theory of adjudication. The formalism or "mechanical jurisprudence" of the day held,[39] roughly, that judges decide cases by legal rules and reasons.[40] By contrast, the realists held that legal rules and reasons were window-dressing for judicial decision-making. Judges actually decide cases based on the facts of the case. As Leiter (2010) puts it: "judges are largely 'fact-responsive' rather than 'rule-responsive' in reaching decisions."[41] Realists suggested, moreover, that the law was, at least in "hard" cases, rationally indeterminate: the applicable legal rules or reasons did not uniquely determine a particular outcome.[42]

Although all realists accepted these central tenets, they differed substantially on what determines how judges respond to facts. Leiter (2010) describes the split as two wings: the Idiosyncrasy Wing and the Sociological Wing.[43] Broadly, the Idiosyncrasy

[37] Posner (2016) at 86-87.

[38] In the case of originalism, for example, this could entail promotions for individual judges endorsing originalism. Note, however, that Judge Posner is a prominent exception, given his criticism of originalism despite being a Reagan appointee.

[39] See Pound (1908).

[40] Leiter (2005).

[41] Leiter (2010) at 249; cf. at 257.

[42] See also Leiter (2005) (discussing how realists also subscribed to a *causal indeterminism* thesis).

[43] See Leiter (2010) at 257-262.

Wing (typified by Jerome Frank) held that the individual personality, ideology, psychology, and other background characteristics of judges determined how they respond to facts. By contrast, the Sociological Wing held that certain general social facts determined judicial outcomes. The difference is largely, therefore, one of degree: social scientists today might say the split was in terms of the degree of heterogeneity in the social facts that affect judicial decision-making.

As Leiter (2005) observes, the realists were proponents (in theory, if not in practice) of scientific methods of discovery and empirical testing.[44] Because most realists thought certain general social facts determined judicial outcomes, many were interested in discovering which factual patterns result in certain judicial responses. Some of the realist works then are, unsurprisingly, descriptive empirical projects to specify how particular fact patterns result in certain types of judicial decisions.

Although largely a descriptive theory of adjudication,[45] the realists also impacted the development of "law in action" and the practice of lawyering. For example, realists worked to reform, improve, and summarize the law of contracts. Corbin, for example, wrote his famous treatise on contract law, whereas Karl Llewellyn was instrumental in drafting and promoting the Uniform Commercial Code. Other realists served as judges (Cardozo, Frank, Holmes, Posner) and in executive agencies (Frank, Cohen).

### 3.1.2 Data Collection

Although many scholars undoubtedly contributed to legal realism, we focus here on the more foundational and influential realist works. Specifically, we examine 16 prominent legal realists and their major works.[46] The 16 writers and their major works are included in Table 1.[47] For each work, we then created two annual time series: (1) academic citations, and (2) judicial mentions.

[44] See Leiter (2005) at 50-51.

[45] See Leiter (2010).

[46] These works were chosen by using the primary sources in survey pieces on American Legal Realism. See, e.g., Leiter (2005) and Leiter (2010). Also, note that although Holmes' *The Common Law*, which was mentioned above in the Posner quote was surely influential, the name of the work is too common to obtain a reliable academic citation series.

[47] We note that many of these works are often considered to be canon. See, e.g., Kennedy and Fisher (2006).

**Table 1.** Legal Realists and Their Works

| Writer | Work |
|---|---|
| Arthur Corbin | Corbin on Contracts |
| Benjamin Cardozo | The Nature of the Judicial Process |
| Benjamin Cardozo | The Paradoxes of Legal Science |
| Felix Cohen | Transcendental Nonsense and the Functional Approach |
| Herman Oliphant | A Return to Stare Decisis |
| Jerome Frank | Are Judges Human? |
| Jerome Frank | Law and the Modern Mind |
| John Chipman Gray | The Nature and Sources of the Law |
| Karl Llewellyn | A Realistic Jurisprudence |
| Karl Llewellyn | Remarks on the Theory of Appellate Decision and the Rules or Canons about How Statutes are to be Construed |
| Karl Llewellyn | Some Realism about Realism |
| Karl Llewellyn | The Bramble Bush |
| Karl Llewellyn | The Common Law Tradition: Deciding Appeals |
| Leon Green | The Duty Problem in Negligence Cases |
| Leon Green | The Judicial Process in Torts Cases |
| Max Radin | In Defense of an Unsystematic Science of Law |
| Max Radin | Law as Logic and Experience |
| Max Radin | Statutory Interpretation |
| Max Radin | The Theory of Judicial Decision: Or How Judges Think |
| Morris Cohen | Property and Sovereignty |
| Morris Cohen | The Basis of Contract |
| Oliver Wendell Holmes | The Path of the Law |
| Richard Posner | The Problematics of Moral and Legal Theory |
| Robert Hale | Force and the State |
| Roscoe Pound | Mechanical Jurisprudence |
| Roscoe Pound | The Scope and Purpose of Sociological Jurisprudence |
| Underhill Moore | An Institutional Approach to the Law of Commercial Banking |
| Underhill Moore | Law and Learning Theory: A Study in Legal Control |
| Underhill Moore | Legal and Institutional Methods Applied to the Debiting of Direct Discounts |
| Underhill Moore | Rational Basis of Legal Institutions |
| Wesley Hohfeld | Some Fundamental Legal Conceptions as Applied in Judicial Reasoning |

The first time series is academic citations of the realist work. For each work, we use data from the JSTOR Data for Research (DfR) program to generate a citation series. This data contains information on the articles that cite the author's last name and the title of the work. The outcome citations are from academic articles only and not, e.g., popular coverage of the work. Because some of the works have very general titles (e.g. "Statutory Interpretation"), the search we carried out is for the author's last name and an exact match to the title of the realist work. This is to avoid too many false positives, i.e. matches that do not actually mention the realist work.

The second time series is how many times courts mention realist works, which serves as our "treatment" path data.[48] To collect judicial mentions, we use public data from CourtListener. This data is then used to encode treatment paths for each work, which represent the number of times a particular work is mentioned each year. For example, the treatment path $\vec{k} = (0, 0, 1, 1, 2)$ means a work receives zero court mentions in years 1 and 2, one court mention in years 3 and 4, and two court mentions in the final year. To avoid false positives in judicial mentions, the search we carried out was for the author's last name and an exact match to the title of the realist work.

We show the cumulative counts aggregated across realist authors in Figure 1. The left panel is the cumulative JSTOR citations series over time. This is the cumulative number of works on JSTOR that cite to one of the Table 1 realist works. The right panel is the number of court mentions. This is the cumulative number of published judicial opinions that mention one of the realist works.

[48] See also below Section 4.2.

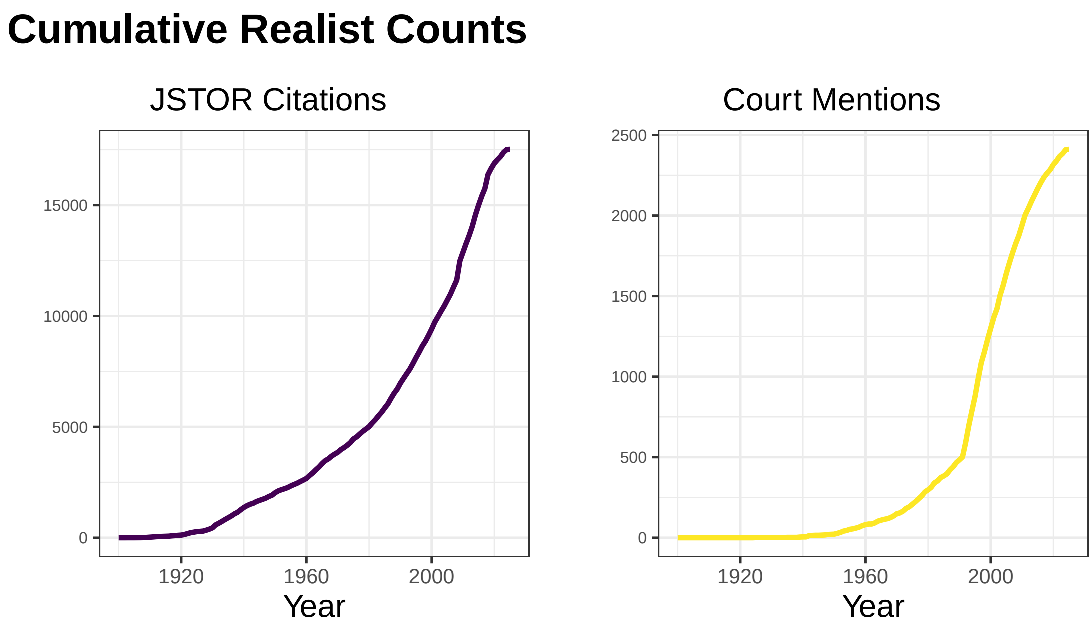


**Figure 1.** Cumulative Realist Counts

*Note:* The $y$-axis differs by panel. Academic citations are more numerous than court mentions for every year.

In the appendices, we also show plots for individual authors. In Figure A1, for example, we plot the cumulative citation series for each work separately. In each author panel, the citation series are shown for the works in Table 1. We also show the citations on a yearly, non-cumulative basis in Figure A2. In Figure A3, we show the yearly court mentions for each author (panels) and work (lines). For example, in 1996 Arthur Corbin's *Corbin on Contracts* was mentioned in 79 court cases. By contrast, Holmes' "The Path of the Law" and Cardozo's "The Nature of the Judicial Process" were each mentioned 8 times by courts in 1996.

The works are mentioned at different times and levels. In Figure 2, we show that some works are highly mentioned early on relative to their publication date (e.g. *Corbin on Contracts*). By contrast, some works only receive judicial notice many years after initial publication (e.g. Radin's "Statutory Interpretation") or are never mentioned by the opinions available on CourtListener (e.g. the works of Underhill Moore). The fact that so few works share similar treatment paths (i.e. treated at the same time and level) is important to the method we use to identify causal effects. Because our method can be used on individual works,[49], the fact that there is little overlap in treatment paths requires modifications to simple causal inference strategies.

[49] As in, e.g., the synthetic control literature pioneered by Abadie et al. (2010).

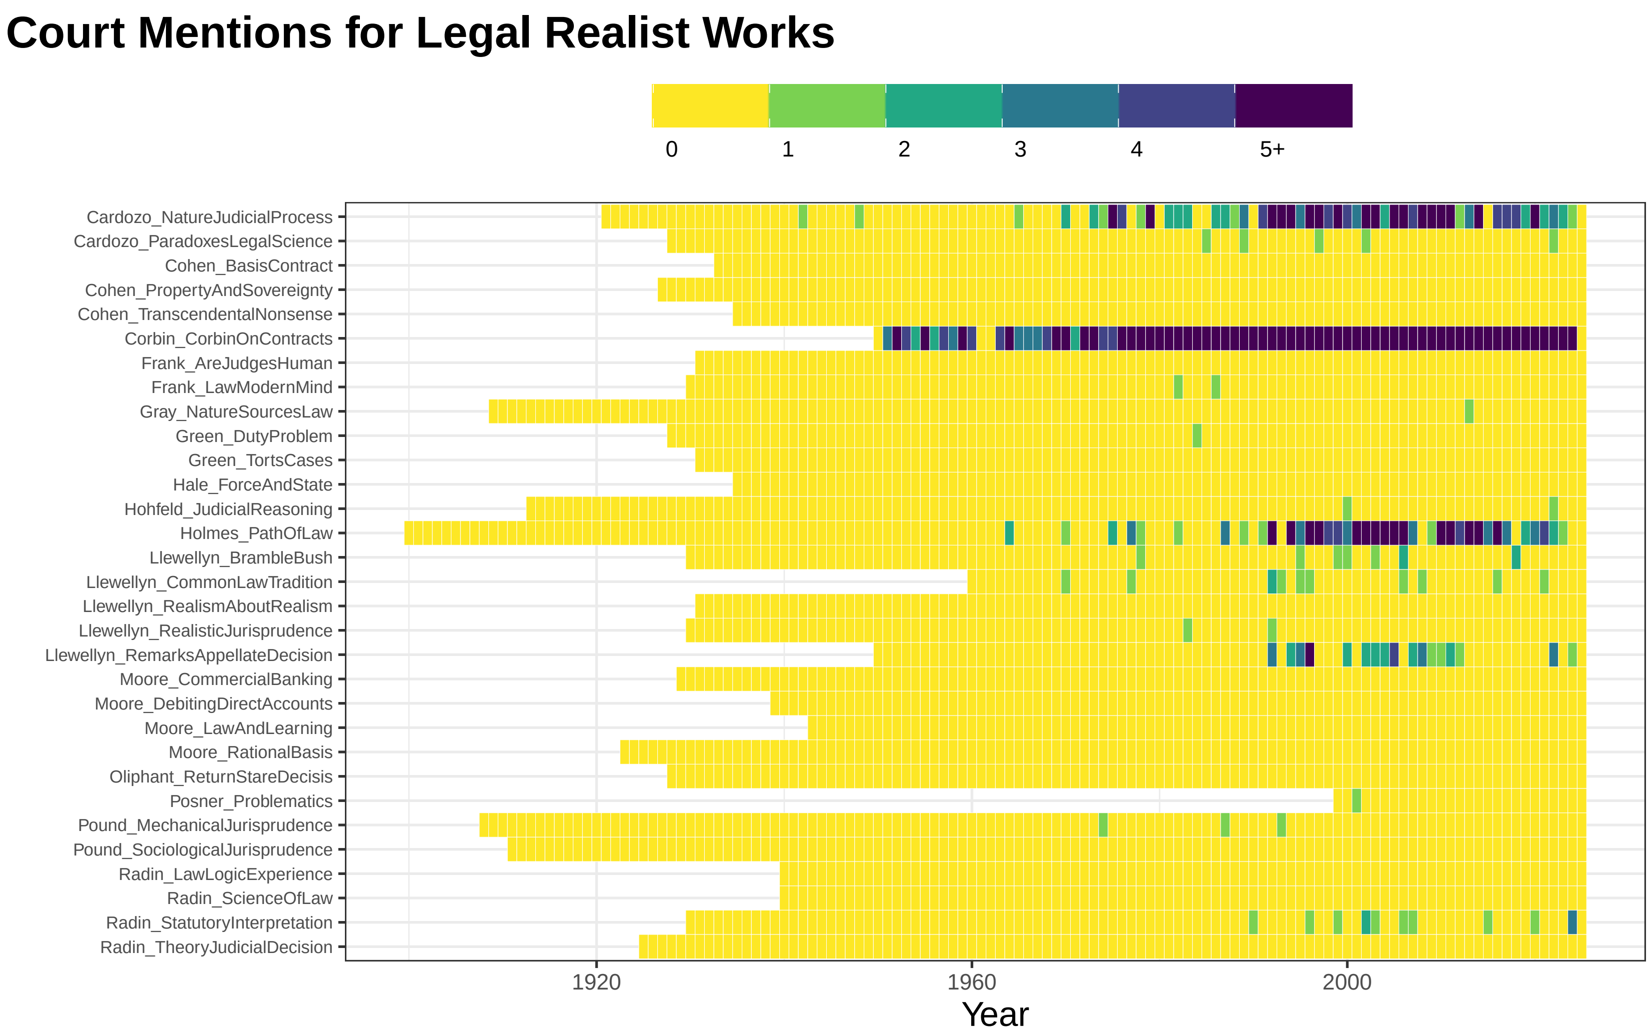


**Figure 2.** Legal Realists: Court Mentions (Treatment Paths)

### 3.2 Originalists

#### 3.2.1 Description of the Movement

Originalism is one of the most influential legal movements of the modern era. The most basic tenet of originalism is that the authoritative meaning of constitutional text is the meaning at the time of enactment.[50] This proposition, originalism proponents argue, restrains judges, restricts the permissible readings of ambiguous constitutional texts, and is "our law."[51]

Originalism is unlike legal realism in at least two important ways. First, it is concerned exclusively with constitutional law. The realists, by contrast, studied across legal subjects and were particularly interested (as academic and practitioners) in commercial law.[52] Second, originalism is both a descriptive and normative theory of law. Whereas realism was largely a descriptive theory of adjudication, originalism, by contrast, is primarily a normative or prescriptive theory for how judges should decide constitutional law issues.[53] Accordingly, many originalist works advance particular "originalist readings" of a constitutional clause and argue that is how judges should decide cases involving that clause.

Although originalism began as a contested descriptive and normative academic theory, it's influence on the judiciary is now widely appreciated. Originalist theory has impacted how judges justify their decisions in constitutional cases.[54] And the rise in judicial acceptance of originalism has generated more demand for originalist scholarship. This is an almost uniquely American phenomenon: originalism in constitutional theory is not widely embraced outside the United States.[55]

#### 3.2.2 Data Collection

Unlike legal realism, which was largely associated with two law schools (Columbia and Yale),[56] originalism claims many varied adherents across law schools with methodological and substantive differences. Because it is also an *active and ongoing* legal

[50] See, e.g., Whittington (2004).

[51] Baude (2015)

[52] See above for some discussion of realism and contract law.

[53] But note that a few originalists also argue for the descriptive claim that originalism is the law. See Baude (2015).

[54] For example in *District of Columbia v. Heller*, both the majority and dissenting opinions devote countless pages to uncovering the 18th century meaning of the Second Amendment.

[55] See Greene (2009a) and Greene (2009b).

[56] Leiter (2005) at 51.

movement, the foundational originalist authors and texts are not as clear as the legal realism case. Accordingly, we have included many more possible foundational originalist works than in the American Legal Realism case.

To construct our population of originalist texts, we used a syllabus for an originalist class ("Originalism and Its Discontents") at the Harvard Law School taught by Professor Stephen Sachs.[57] We used this syllabus because it is relatively exhaustive—it contains a long list of optional readings relevant to originalism—and, therefore, allowed us to cast a wide net in our search for "foundational" originalist texts. We omitted all non-academic articles (e.g. court cases) that are listed on the syllabus, as well as works for which we could not find academic citation series in JSTOR or for which the titles of the work were so common (e.g. "Commerce Clause") that they would not reflect citations to the academic article alone.

Our final list contains 191 originalism-relevant texts. As in the realist case, the list of works is neither exhaustive nor definitive. Nevertheless, it includes many popular and influential originalist texts. Importantly, note that not all of the works are by authors that adhere to originalism. Some of the works are merely *relevant* to originalism because this is a course syllabus.

As before, we generated two time series for each originalist work. The first is the number of academic citations to each work (the JSTOR citation series). The second series is how many opinions each originalist work is mentioned in each year. The cumulative time series are plotted in Figure 3: the cumulative JSTOR citations series aggregated across authors is displayed in the left panel; the cumulative court mentions series aggregated across authors is displayed in the right panel.

[57] We used the Spring 2025 Sachs syllabus available online at https://www.stevesachs.com/syllabi/originalism.pdf.

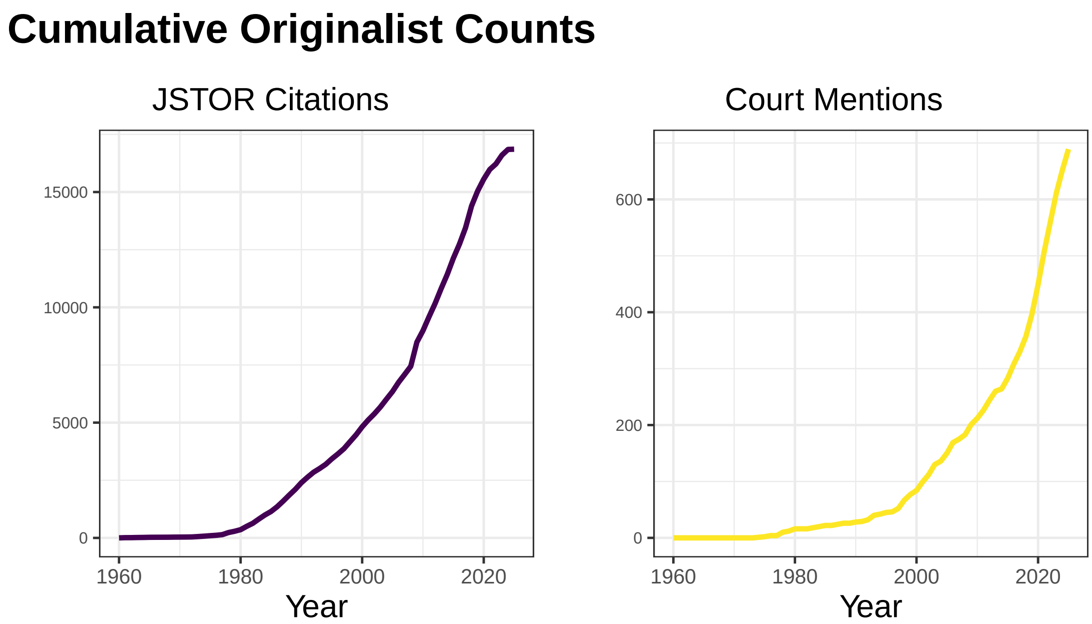


**Figure 3.** Cumulative Originalist Counts

*Note:* The $y$-axis differs by panel.

As in the realism case, the originalist works are mentioned by judges to varying degrees and at varying times. The timelines in Figure C1, Figure C2, Figure C3, and Figure C4 show the levels and years the different works are mentioned in published judicial opinions. As above, it is important to notice here that the "treated" works experience little overlap in the treatment path space.[58] Unlike the realist case, however, there are many works that are never treated (i.e. never mentioned in a judicial opinion). This is not unexpected because originalism is a newer legal theory than realism.

## 3.3 Descriptive Findings

### 3.3.1 Textual Data

With the data collected, we are primarily interested in the academic citation and judicial mention metrics. However, we note here briefly that the JSTOR data in particular contains interesting textual data as well. For each realist or originalist work, the JSTOR data includes matches to articles or books that cite that work. Each match contains information on the keywords used in the article that cite the

[58] That is, they do not experience the same number of judicial mentions at the same times. Because of this, we group works into simplified treatment cohorts. See Section 4.2 below.

realist or originalist work. For example, in the JSTOR data for Jerome Frank's *Law and the Modern Mind*, the first article returned that cites Frank's work is Rumble (1965). The JSTOR data includes the following keyphrases for the Rumble article:

> legal realists; legal realism; llewellyn; karl llewellyn; jurisprudence; judicial; american legal; legal norms; ideals; judicial decisions

Each article that cites one of the realist or originalist works has its own keyphrases. In Figure 4, we show the 100 most common word stems among these keyphrases for both the realist corpus and the originalist corpus.

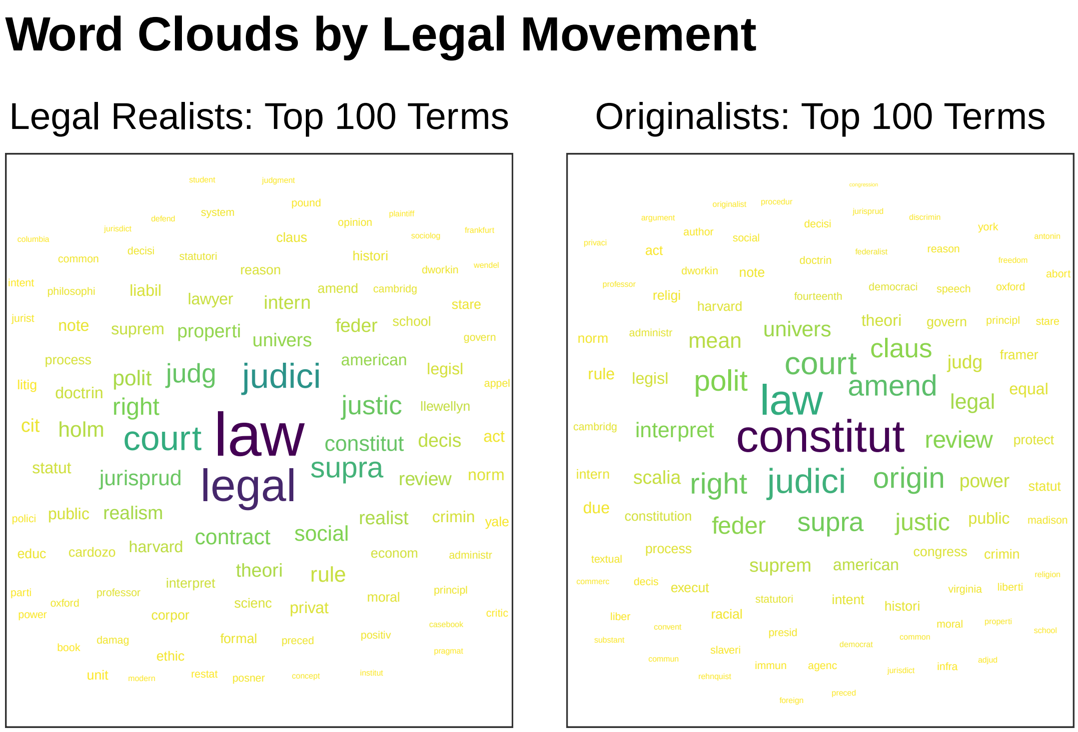


**Figure 4.** Word Clouds by Legal Movement

The two word clouds contain 54 words in common. These are mostly words that relate to law and judging generally. The three most mentioned words in articles citing legal realist works that are also mentioned by originalist-citing articles are: law, legal, and judici. The three most mentioned word stems in articles citing originalist works that are also mentioned by realist-citing articles are: constitut, law, and judici.

Unsurprisingly, the academic articles that cite our originalist corpus more frequently have keyphrases involving constitutional issues and particular constitutional clause. The three most mentioned words in articles citing legal realist works that are *not* mentioned by originalist-citing articles are: contract, realism, realist. The three most mentioned words in articles citing originalist works that are *not* mentioned by realist-citing articles are: origin, mean, and scalia. The originalist-citing articles also include many keyphrases realted to constitutional clauses that are not included in the top 100 realist keyphrases, such as: fourteenth, equal, protect, religion, and abort.

Although it is beyond the scope of the present article, scholars might be able to identify more works that may reasonably be called originalist or realist works by using textual data. The keyphrases (or the text of the originalist and realist works themselves) could be used to train an LLM classifier to categorize academic articles into particular legal-theoretical movements.

#### 3.3.2 Citations & Mentions

The primary data of interest are the two time series for each work. Recall the two time series are: (1) the number of academic citations to an originalist or realist work (JSTOR Citations); and (2) the number of mentions for an originalist or realist work in published judicial opinions (Court Mentions).

In Figure 5, we show the basic linear fit for each author's total JSTOR citations on the total court mentions. The totals are plotted on the log scale for visualization purposes. We label the names for authors with four or more works in each dataset. The linear fits for the realists (yellow line) and originalists (purple line) both have positive slopes. That is, the total number of academic JSTOR citations is increasing in the total number of court mentions.

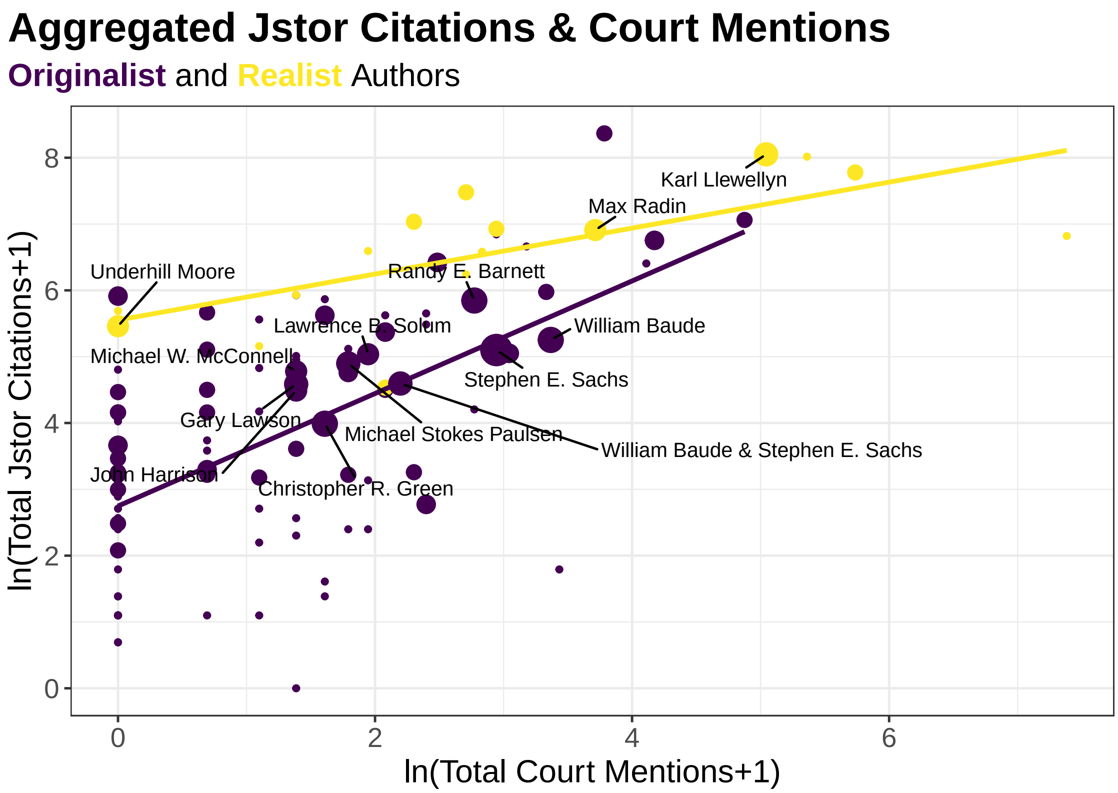


**Figure 5.** Overview: Relationship between Court Mentions and Academic Citations
*Note:* Each dot represents a unique author grouping and has volume proportional to the number of total works. Authors with four or more total works are labeled.

Estimates calculated using regressions, an example of which is the line of best fit displayed above, does not necessarily imply a causal relationship. Indeed, this is a complicated setting with time-varying treatments (court mentions) where a causal relationship, in so far as it exists, manifests with some time lag.

The regression specification that produced the line of best fit above can be spelled out as follows:

$$\text{Total JSTOR Citations}_{ig} = \alpha_g + \beta_g \times \text{Total Court Mentions}_{ig} + \text{error}_{ig}$$

The subscripts $g$ relate to the author group (realists or originalists). The independent variable Total Court Mentions$_{ig}$ is the total court mentions for author $i$ in group $g$.[59] The goal is to obtain an estimate for $\beta_g$ (the slope in the line of best fit). Such an estimate can be interpreted as the average increase in citations that result from one

[59] Note also that the variables are on the log scale in the specifications.

more court mention for the author group.[60]

Since this specification obscures important variation within publications in the author group and variation in the court mention path, a different regression that is at a more granular work and year level can be also specified. Now the $i$ subscript refers to the published work, the subscript $t$ to the year, and the subscript $g$ to the group of authors (realists or originalists):

$$\text{Cumulative JSTOR Citations}_{igt} = \alpha_g + \beta_g \times \text{Cumulative Court Mentions}_{igt} + \text{error}_{igt}$$

All citations and mentions are calculated as cumulative series as in Figure 1 and Figure 3. Now the interpretation of $\beta_g$ is the average increase in academic citations in a year for an additional court mention in that year (for group $g$).

Notice, however, that this specification assumes that an effect is instantaneous. Given the possible mechanisms by which judges learn of and mention legal scholarly works in their opinions, such an assumption is overly restrictive. As Figure 2 illustrates, scholarly works are never mentioned immediately after publication. Therefore such an instantaneity assumption can induce a downward bias on the estimate $\beta_g$ for each group of authors $g$.

There are alternatives to account for the violation of instantaneity. But they require either culling the sample or making strong assumptions about the judicial process. For example, we may consider works only after their first judicial mention. Or we may assume that works take some number of years to be recognized by the judiciary after publication, and truncate both series to account for this lag. Both are imperfect solutions and likely to overestimate the effect of judicial acknowledgments because the information removed occurs where judicial mentions are small.

Regression evidence, moreover, may be particularly problematic for causal inference where, as here, there is sparsity in the treatments. That is, we only observe limited court mention paths out of a large space of conceivable such paths. Estimates such as $\beta$ that attempt to encode this process then have limited information to "chew on". To make matters worse, naive linear best fit regression makes comparisons between units on different causal paths and does not necessarily have an obvious causal interpretation.

The previous discussion implies that greater care is needed to spell out the coun-

[60] While the linear fit seems appropriate to the realists, it cannot account for a cluster of originalist works that receive no court mentions but have academic citations that exceed author groups that received many court mentions.

terfactual assumptions necessary to make causal conclusions while relaxing instantaneity. We now turn to detailing the challenges to causality in this setting. We then spell out the assumptions needed to bound the causal effect of court mentions on academic citations.

## 4 How to (Partially) Identify Causal Effects

The canonical machinery of causal inference is derived under simple treatment paths. The familiar potential outcomes framework is introduced in a setting where each unit's treatment path is only one time period long. There are then only two possible treatment paths: treated or untreated. It is assumed there are a sufficient number of treated and untreated units that can be used to impute the counterfactuals for each potential outcome. Under different assumptions, quantities like the average treatment effect, average treatment effect on the treated, or average treatment effect on the untreated can be estimated.

In dynamic treatment settings, treatments paths are many-periods long. The potential outcomes framework still applies but is more involved. More assumptions must be made, for example, about whether past treatments matter for later treatments, whether units dynamically select treatments (e.g. "dynamic selection on gains"), and whether units anticipate treatment. Yet, even in these dynamic settings, some treatment paths are simpler than others. If all treatment occurs at one period, the paths can essentially be reduced to the simpler setting by making simplifying assumptions.[61]

Because the number of potential outcomes grows for each individual as treatment paths increase in complexity, we now turn to certain simplifying assumptions on how full treatment history affects potential outcomes, which reduces but does not eliminate the complexity. We then propose our Bounded Deviations strategy and compare it to the canonical difference-in-differences (DID) approach. We show how these assumptions can point-identify or bound causal effects of interest in our setting.[62]

This section first describes some intuition for why causal inference is difficult in a setting where outcomes can depend on the entire treatment history. Then we set up the mathematical framework for this setting and describe how it creates a

[61] A well-studied dynamic setting is the staggered adoption setting where the treatment may occur at different times of the path, but any unit if treated at all is only treated once. In such a setting, treated units are usually assumed to retain memory of treatment forever. See, e.g., Callaway and Sant'Anna (2021). See also discussion in Appendix B.

[62] Our approach is a Manski-type partial identification assumption. See Manski and Pepper (2018) for further details.

dimensionality problem. Finally, we briefly discuss how some simplifying assumptions can be difficult to motivate, hard to interpret, and result in biased causal estimates.

### 4.1 Counterfactual Framework

We now introduce notation to describe the potential outcomes as a function of the treatment histories or treatment paths for each individual $i$ in each time period $t$. Consider a setting with time periods $t \in \mathcal{T} = \{1, \ldots, T\}$ and let $\vec{K}_i = (k_{i1}, \ldots, k_{iT})$ be the history of observed treatments for individual $i$ at each time period. Each $k_{it}$ is the treatment level $i$ receives at time period $t$. Note that, in principle, each $k_{it}$ can be multidimensional if there are multiple treatments, but we focus on scalar $k_{it}$ here. Suppose further that the treatments are discrete: for each individual $i$ and time period $t$, $k_{it} \in \mathcal{K} = \{0, 1, \ldots, K\}$.

We write the potential outcomes for each individual $i$ in time period $t$ under treatment path $\vec{k} = (k_1, \ldots, k_T)$ as

$$Y_{it}(\vec{k}) = Y_{it}(k_1, \ldots, k_T)$$

Implicit in this notation is that there are no treatment "spillover" effects, i.e. the potential outcome for unit $i$ at time $t$ depends only on $i$'s treatment path. Finally, we define the potential outcome matrix $\boldsymbol{Y_i}$ for individual $i$. For a fixed ordering of the potential treatment paths $\mathcal{K} = \{\vec{k}_1, \vec{k}_2, \ldots\}$, the $(i, j)$-th entry of $\boldsymbol{Y_i}$ is given by $Y_{i,t=j}(\vec{k}_i)$. For example, consider a two-time period model with a binary treatment possible in each period and an individual $i$ with observed treatment path $\vec{K}_i = (k_{i1} = 0, k_{i2} = 1)$. Suppose further the treatment paths are ordered as follows: $\mathcal{K} = \{(0,0), (0,1), (1,0), (1,1)\}$. Then we write $i$'s potential outcome matrix $\boldsymbol{Y_i}$ as follows:

$$\boldsymbol{Y_i} = \begin{bmatrix} Y_{i,t=1}(0,0) & Y_{i,t=2}(0,0) \\ Y_{i,t=1}(0,1) & Y_{i,t=2}(0,1) \\ Y_{i,t=1}(1,0) & Y_{i,t=2}(1,0) \\ Y_{i,t=1}(1,1) & Y_{i,t=2}(1,1) \end{bmatrix} = \begin{bmatrix} ? & ? \\ Y_{i,t=1}(0,1) & Y_{i,t=2}(0,1) \\ ? & ? \\ ? & ? \end{bmatrix}$$

The unobservable potential outcomes for $i$ (the question marks) are the **identification problem**: because unit $i$ has treatment history $\vec{K}_i = (0, 1)$, we cannot observe the outcomes for $i$ in period $t$ at different treatment histories.

### 4.2 Curse of Dimensionality & Treatment Path Cohorts

In general, the number of unobservable potential outcomes, grows as the cardinality of $\mathcal{K}$ and $\mathcal{T}$ grow. This creates a curse of dimensionality problem: allowing for unit $i$'s potential outcome at time $t$ to depend on the full treatment path increases the complexity of the counterfactual identification problem.[63]

To take control of the explosion of paths, we define four treatment cohorts, which simplifies the treatment path space. The first cohort is the never-treated cohort: these units never receive any judicial mentions. The remaining three cohorts are those that receive at least one judicial mention. They are split according to whether they receive a court mention Early, Mid, or Late. The years were determined to ensure a roughly equal split among the three treated cohorts. This resulted in the following cohort splits for each legal movement:

**Realists**

| Cohort | Number Works |
|---|---|
| Early (1925–1945) | 9 |
| Late (1976–2025) | 7 |
| Mid (1946–1975) | 7 |
| Never-Treated | 8 |

**Originalists**

| Cohort | Number Works |
|---|---|
| Early (1974-2009) | 31 |
| Late (2020-2025) | 32 |
| Mid (2009-2019) | 28 |
| Never-Treated | 100 |

**Table 2.** Cohort Distribution for Realists and Originalists

In each table the years in parentheses indicate the times that the cohort first received a court mention. For example, if a realist work was first mentioned in a published court opinion between 1925–1945, then it is in the Early treated cohort. The cohorts are roughly equal in terms of size, except that for the originalists there are many more Never-Treated units.

### 4.3 Bounded Deviations & Partial Identification

We now turn to explaining the framework we use to bound treatment effects in this setting. Specifically, our framework is a variant of the bounded variation framework discussed in Manski and Pepper (2018). As in Manski and Pepper (2018), we propose assumptions that to bound (i.e. partially identify) causal effects for a given tuning

[63] To address the dimensionality problem, researchers typically restrict the dependence of the potential outcomes $Y_{it}(\cdot)$ on the treatment history. We discuss further in Appendix B how different assumptions that are (implicitly or explicitly) used in practice can reduce the dimensionality of the identification problem.

parameter. Our approach—which we call **Bounded Deviations**—uses information from units before their treatment paths diverge (i.e. *pre-divergence*) to impute counterfactual paths when the treatment paths no longer agree (i.e. *post-divergence*). Intuitively, our method imposes that the average treated group counterfactual must be "close" enough to the average control group observed outcome.[64]

The framework is straightforward to apply, interpret, and estimate. We also compare our Bounded Deviations method to the widespread difference-in-differences method (DID), where we discuss how the bounds from our method contain the usual DID estimates if parallel trends does hold. We also illustrate how causal conclusions are sensitive to the tuning parameter assumption.

To begin, we specify the treatment effect estimand of interest. We study the *average treatment effect on the treated* ("ATT") for units on two treatment paths, which we define as follows: consider for treatment paths $\vec{k}$ and $\vec{k}'$ the average treatment effect in time $t$ on units treated with history $\vec{k}$ as:

$$ATT(t, \vec{k}, \vec{k}') = \mathbb{E}[Y_{it}(\vec{k}) - Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}]$$

This is the average treatment effect from switching to path $\vec{k}'$ for units on path $\vec{k}$ at time $t$. As discussed above, the identification problem is that $Y_{it}(\vec{k}')$ is unobservable for units with observed treatment history $\vec{K}_i = \vec{k}$.

Consider two distinct treatment paths $\vec{k}$ and $\vec{k}'$, and let the first time period the treatment histories do not agree be denoted $T_{div} = \text{argmin}_t(k_t \neq k'_t)$. We use information from the pre-divergence periods (i.e. time periods before $T_{div}$) and the following assumption to impute counterfactual outcomes.

**Assumption 1.** *[Average Bounded Deviation] Let $T_{div} = \arg\min_t(k_t \neq k'_t)$ for distinct treatment paths $\vec{k}$ and $\vec{k}'$. Then the following holds for each $i$, $t \geq T_{div}$, any $(\vec{k}, \vec{k}')$ pair, and some $C \in \mathbb{R}_+$:*

$$\left|\mathbb{E}[Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}']\right| \leq C \cdot \max_{t<T_{div}} \left|\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}']\right|$$

*where $C = C(t, \vec{k}, \vec{k}')$ can depend on the time period and both treatment paths.*

This assumption says the counterfactual average outcome from treating with $\vec{k}'$ instead of $\vec{k}$ for units with observed treatment path $\vec{K}_i = \vec{k}$ is "near" the observed

[64]Specifically, our method requires the treated group counterfactual to be within the maximum pre-path divergence difference in observed outcomes of the observed control average outcomes. See below.

outcome for units on treatment path $\vec{k}'$. How near? Specifically, the counterfactual average for the $\vec{k}$-treated group is within $C$ times the the maximum absolute difference in means for units on the two paths before the paths diverge.[65]

This assumption has "identifying power" because it relates the unobservable counterfactual outcomes—i.e. the conditional expectation $\mathbb{E}[Y_{it}(\vec{k}')|K_i = \vec{k}]$—to observable averages. Under this assumption, the average treatment effect on the treated (ATT) is partially identified. That is, we may obtain upper and lower bounds on the ATT. The bounds are given by the following observable quantities for fixed $C$:

$$\overline{ATT(t, \vec{k}, \vec{k}')} = \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \Big(\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}'] - C \cdot \text{Pre-Div}(\vec{k}, \vec{k}')\Big) \tag{1}$$

$$\underline{ATT(t, \vec{k}, \vec{k}')} = \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \Big(\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}'] + C \cdot \text{Pre-Div}(\vec{k}, \vec{k}')\Big) \tag{2}$$

where the maximum pre-divergence deviation is given by:

$$\text{Pre-Div}(\vec{k}, \vec{k}') = \max_{t<T_{div}} \Big|\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}']\Big|$$

In this notation, the upper bound is given by $\overline{ATT(t, \vec{k}, \vec{k}')}$ and the lower bound is given by $\underline{ATT(t, \vec{k}, \vec{k}')}$.

We illustrate the Bounded Deviations assumption and its identifying power below. In Figure 6, we consider two conditional expectation functions on different treatment paths $\vec{k}_1$ and $\vec{k}_2$. For simplicity, we shall call the $\vec{k}_1$ path "Treated" and the $\vec{k}_2$ path "Untreated." These treatment paths are the same until 1950, when they diverge. The observed conditional expectation $\mathbb{E}[Y_{it}(\vec{k}_1|K_i = \vec{k}_1]$ is given by the blue line from 1900 to 2000. We show the observed conditional expectation $\mathbb{E}[Y_{it}(\vec{k}_2|K_i = \vec{k}_2]$ for units on path $\vec{k}_2$ in red only before 1950 for visual clarity. The maximum deviation between the average for the two groups is 1.69 and occurs shortly after 1925.[66]

The problem of causal inference here (i.e. the "identification problem") is what counterfactual outcome would units on one path experience if they were instead on

[65] Note that $C = C(t, \vec{k}, \vec{k}')$ can depend on the time period $t$ and treatment paths ($\vec{k}$, $\vec{k}'$). In what follows, we consider fixed $C$-values that are constant across all post-divergence time periods. But, in principle, further assumptions on $C$ as a function of the treatment paths $\vec{k}, \vec{k}'$ or the time period can be incorporated. For example, in some cases there may be good reasons (e.g. theory, past studies, etc.) to believe the bound should grow in time. That is, $C$ should be a weakly increasing function of $t$.

[66] Note that what is plotted in Figure 6 are the conditional expectation functions (CEFs) for units on a particular path. For example, the solid-red line is the average observed outcome for units on path $\vec{k}_2$ in each time period.

the other treatment path? In the left panel of Figure 6, the blue-striped region is the counterfactual region for the blue units (i.e. the units that actually experienced treatment path $\vec{k}_1$) if they had instead experienced treatment path $\vec{k}_2$. Recall our identifying assumption: if the two treatment paths had not diverged, units on treatment path $\vec{k}_1$ would be "near" the observed outcome for units on path $\vec{k}_2$. How close would they be? They'd be within the maximum pre-divergence deviation.[67] In Figure 6, the maximum pre-divergence deviation is 1.69 and the blue-striped region is the counterfactual region under the assumption that blue units would've continued to be within 1.69 of the observed red unit outcomes.

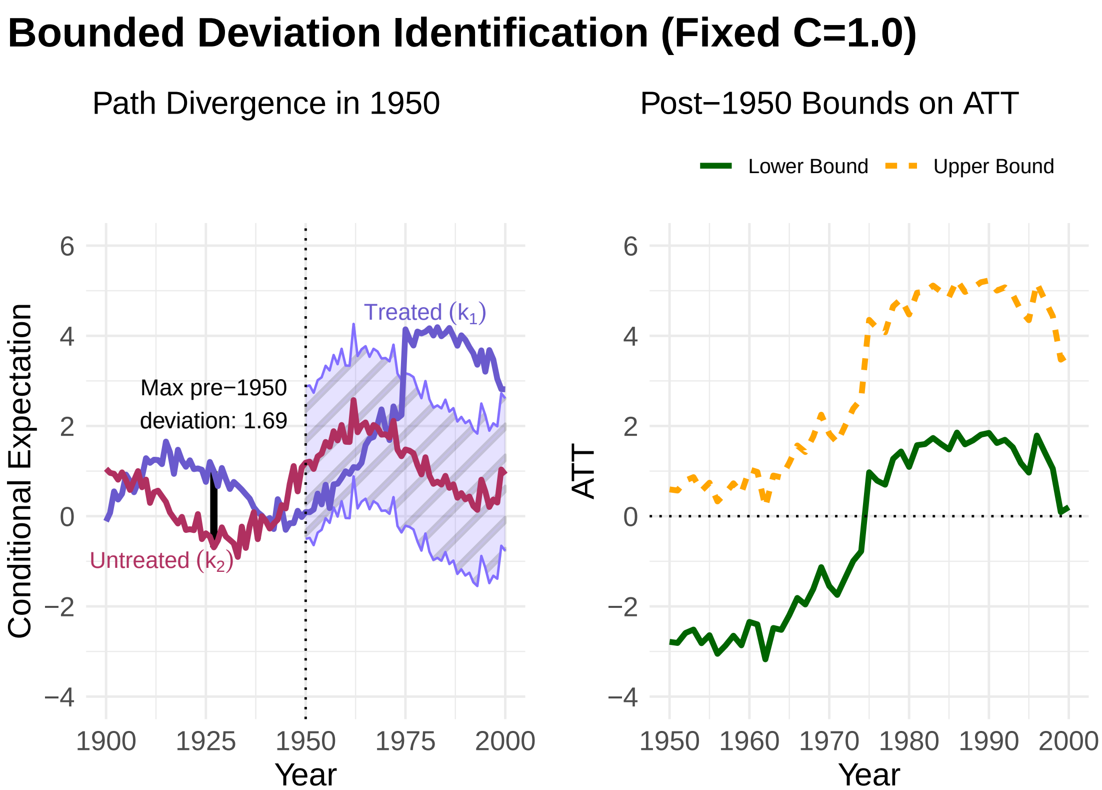


**Figure 6.** Bounded Deviations (Fixed $C = 1$)

*Note:* In the left panel, the two lines are $\mathbb{E}[Y(\vec{k})|\vec{K}_i = \vec{k}]$ for two treatment paths: $\vec{k}_1$ (blue) and $\vec{k}_2$ (red). The black dotted vertical line is the path divergence time period ($T_{div}$). The blue-striped region is the counterfactual region for blue units (i.e. on treatment path $\vec{k}_1$) if they had instead experienced the red unit treatment path ($\vec{k}_2$). The maximum pre-divergence deviation is 1.69 and is shown by the black vertical line around 1925. In the right panel, $y$-axis is $ATT(t, \vec{k}_1, \vec{k}_2) = \mathbb{E}[Y_{it}(\vec{k}_1) - Y_{it}(\vec{k}_2)|\vec{K}_i = \vec{k}_1]$. The black dotted horizontal line is plotted for $ATT = 0$. The upper and lower bounds on the $ATT$ for each post-1950 time are plotted as dashed-yellow and solid-green lines, resp.

[67] When the multiple $C = 1$.

The bounds on the treatment effects are also shown in Figure 6. In the right panel, the lower bound is plotted as a solid-green line and the upper bound is plotted as the dashed-yellow line. The dotted horizontal line at zero is also shown and corresponds to $ATT = 0$.

What do these bounds tell us? Before 1975, our bounds include the possibility of a zero effect—i.e. the dotted line at zero lies between the lower bound (solid-green line) and the upper bound (dashed-yellow line). Notice that this corresponds to the pre-1975 solid-blue line lying entirely within the blue-striped counterfactual region in the left panel. Moreover, we cannot say (identify) what the direction (or sign) of the effect is. In other words, we cannot know whether $\vec{k}_1$-treated units would experienced higher or lower outcomes relative to their counterfactual outcomes had their experienced treatment $\vec{k}_2$.

After 1975, by contrast, the bounds do not contain the zero effect. That is, the dotted line lies below the lower bound on the $ATT$ (solid-green line). Therefore, post-1975 we can identify that the treatment effect of being on treatment path $\vec{k}_1$ rather than path $\vec{k}_2$ increased the outcomes for units on path $\vec{k}_1$. In other words, the treatment effect is sign-identified post-1975 even if we cannot say precisely the magnitude of the effect.

The bounds we obtain are sensitive to the tuning parameter $C$ in our counterfactual assumption. The plots above assume $C = 1$, which intuitively means that the pre-divergence and post-divergence periods are "roughly the same." How alike they are is governed by how large or small $C$ is. Large values of $C$ will increase the size of the counterfactual region and suggest that the past is a relatively weak guide to the future. By contrast, values of $C < 1$ imply that the counterfactual outcome is even closer to the observed outcome than it was pre-treatment divergence. We show how this affects the upper and lower bounds in Appendix D.

For those familiar with the vast economics literature on difference-in-differences (DID), we also compare our Bounded Deviations method to the simpler DID standard. We consider here the simplest possible setting where there are only two groups—the treated and control—that may receive a single dose of a binary treatment in the year 1950.[68] In the left panel of Figure 7, we show the standard DID strategy to identify causal effects in this simple setting. The solid-blue line is the observed average outcome in each year for treated units. The solid-red line is the observed

[68]For example, this could be: "Did Congress ban the sale of good $i$ in 1950?" If the answer is yes, then good $i$ is part of the treated group. If the answer is no, then good $i$ is part of the control group.

average outcome for the control group for each year.

How to estimate the causal effect of treatment on the treated group? To estimate the treatment effect, the DID-strategy assumes that the counterfactual potential outcome would've trended in parallel to the control group. This is called the *parallel trends* assumption and it is the basis for a vast empirical economics literature.[69] In Figure 7, the dashed-blue line shows this counterfactual, unobserved potential outcome for units that actually received treatment in the left panel. This is the outcome that treated units would have experienced (on average) if they had not been treated. Notice that the dashed-blue and solid-red lines are parallel.

The treatment effect on the treated is the difference between the solid-blue and dashed-blue lines. The difference is the difference between the average treated outcome and the average control outcome for units that actually received treatment. In the left panel of Figure 7, the treatment effect—called the *average treatment effect on the treated* ("ATT")—is shown by the black arrow labeled "Treatment Effect". Notice that in this example, the ATT is constant across the post-1950 years, but this is not a requirement of the DID strategy.

[69] See Goldsmith-Pinkham (2024).

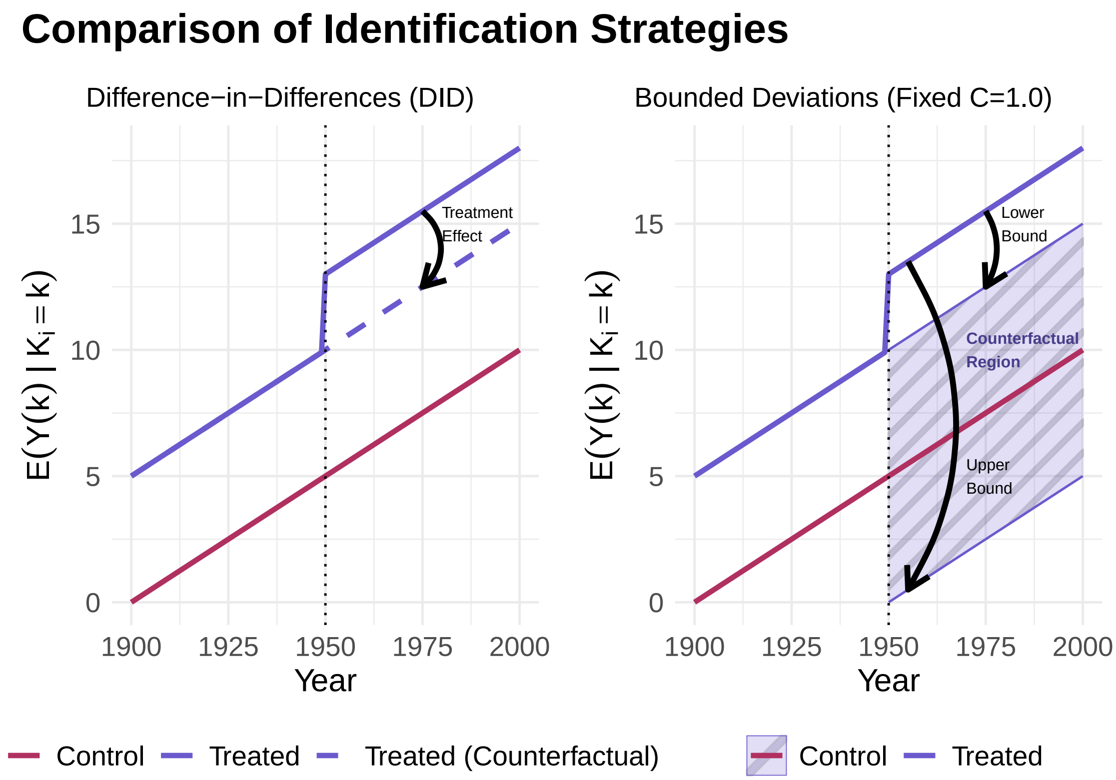


**Figure 7.** Bounded Deviations & DID Comparison

*Note:* The black dotted vertical line is the path divergence time period ($T_{div} = 1950$). The blue-striped region is the counterfactual region.

Our method in this simple setting is depicted in the right panel of Figure 7. The observed average outcomes are, as before, plotted in the solid-blue and solid-red lines. The difference here is that we allow for the counterfactual (i.e. untreated) average outcome for the treated group—which before was the dashed-blue line—to not evolve in parallel with the control group average outcome. Indeed, we allow for the counterfactual average outcome to differ from the parallel counterfactual in the DID provided it remains "close" to the control group outcome.[70] In the right panel of Figure 7, this is shown by the blue-striped "Counterfactual Region". Our method allows the counterfactual (i.e. untreated) average outcome for the treated group to lie above or below the control group average outcome (solid-red line) so long as it remains within the blue-striped region.

The treatment effect on the treated (ATT) under our method is bounded above and below. It is bounded below by taking the difference between the observed average

[70]For much more discussion of how to relax the parallel trends assumption, see, e.g., Rambachan and Roth (2023).

treated outcome and the largest possible counterfactual average outcome. In the right panel of Figure 7, this is shown by the arrow labeled "Lower Bound". The upper bound is found similarly: take the difference between the observed average treated outcome and the smallest possible counterfactual average outcome (labeled "Upper Bound" in Figure 7). Finally, notice that the bounds in our example nest or contain the DID estimated treatment effect. That is, if parallel trends does hold, then the DID-based estimate is equal to the Bounded Deviations lower bound (for $C = 1$).

### 4.4 Estimation & Inference

We also compute all bounds as a function of $C$, which facilitates checking how conclusions change as a function of $C$. The estimation strategy proceeds by first partitioning the units by observed treatment paths. These paths will make up the counterfactual paths of interest. Next, for any two path pair of interest $(\vec{k}, \vec{k}')$ we compute the divergence period $T_{div}$ and the maximum deviation in time periods before $T_{div}$. We then estimate the counterfactual path using the identification assumption for different $C$ values. Although the primary focus of the methodology here is on identification, we note that, where possible, one can estimate the conditional expectations non-parametrically and use bootstrap techniques to compute standard errors. Inference is a difficult problem, however, in cases where there is only one treated unit (as in synthetic control), so we leave further estimation and inference issues for future work.

## 5 Inferring Causation: How Courts Impact Legal Movements

In this section, we use the methods developed above to estimate bounds on the causal effects of judges on two legal movements: legal realism and originalism. We quantify the impact of judicial mentions on the reception of academic legal ideas. Specifically, we bound the causal effects of judicial citation to certain legal realist or originalist works on the popularity of those works over time. We illustrate how to estimate upper and lower bounds for works if they had experienced alternative (i.e. counterfactual) treatment paths.

### 5.1 The Effects of Court Mentions on Legal Realism

In our first application, we estimate the effect of judicial mentions on the reception for the legal realist thinkers mentioned above. Recall that we separated the legal realist works into four treatment cohorts: the Never-Treated cohort received zero court mentions; works in the Early treated cohort first received court mention in 1925–1945; works in the Mid treated cohort first received court mention in 1946–1975; and works in the Late treated cohort first received court mention in 1976–2025.

Our first set of results estimates the effect from being in the Mid treated cohort versus having never been treated. In Figure 8, we plot the upper bound (dashed-yellow) and lower bound (solid-green) on the $ATT$ for being in the Mid treated cohort versus the Never-Treated cohort for units actually in the Mid treated cohort. The different panels are for different values of the identification tuning parameter $C$.[71]

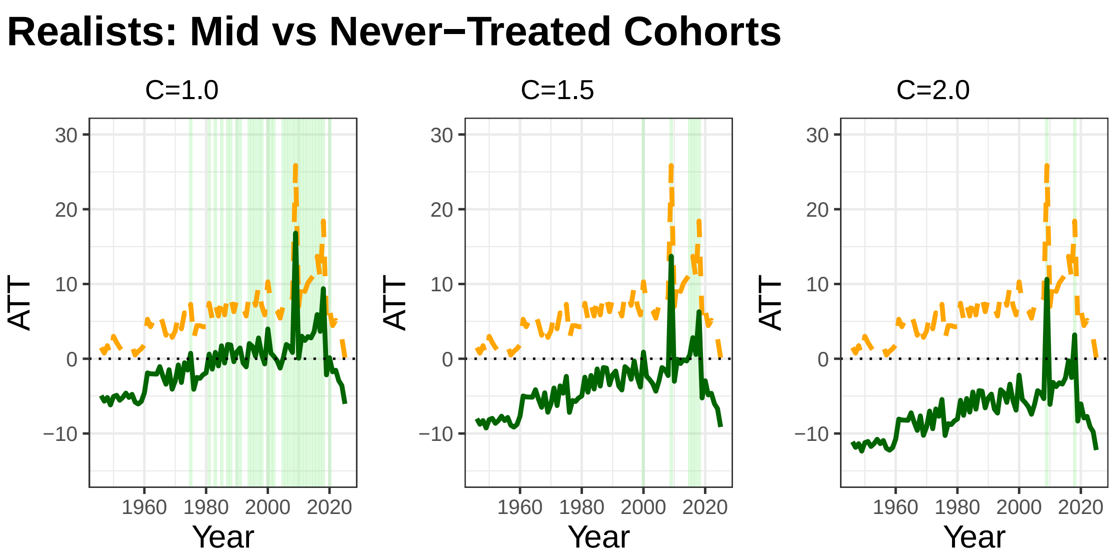


**Figure 8.** Realists: Mid vs Never-Treated Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

The light green-shaded region in each panel are those years for which we identify the sign of the $ATT$ to be positive. As shown in the left-most panel of Figure 8, many of the 1980-2000 years sign-identify the $ATT$ when $C = 1$. This means that the Mid units would have been much "closer" to the Never-Treated units if they had never received judicial mentions. Instead, the judicial mentions increased the popularity (academic JSTOR citations) for realist works that were first treated in the 1946–1975 years (Mid cohort).

The effect is even more pronounced for realist works in the Late treated cohort. Relative to Never-Treated units, realist works that were first mentioned in judicial

[71]See Appendix D for further details.

opinions from 1976–2025 received a positive boost in academic citations. This effect holds across values of the identification tuning parameter $C$, as shown below in Figure 9. In each panel, there are many years for which we sign-identify a positive $ATT$ for Late treated units relative to the Never-Treated cohort. This is indicated by the large number of years which are shaded light green in each panel of Figure 9.

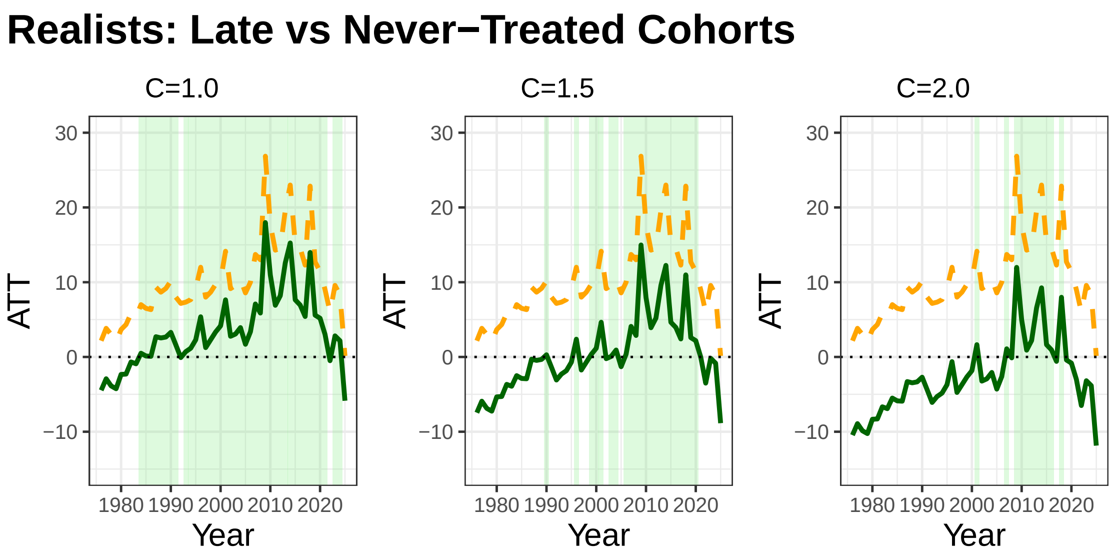


**Figure 9.** Realists: Late vs Never-Treated Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

These findings highlight an important difference between legal scholarship and judicial opinions. Existing empirical studies show that the value of court opinions depreciate rapidly. For example, Black and Spriggs II (2013) find that case precedent is less likely to be cited over time: new cases are about 30% likely to be cited, whereas 20 year old precedent is less than 10% likely to be cited.[72] By contrast, our results suggest some legal scholarship receives more citations later in life.[73] In other words, court opinions age like milk but legal scholarship ages like wine.

Further comparisons between treatment cohorts are explored in Appendix E. To summarize the results there, Early treated units received a positive boost from court mentions relative to Mid treated units, although this effect only appears in the 2000-

[72]See Figure 1 in Black and Spriggs II (2013), which shows that the predicted probability of citation as a function of the age of precedent decreases rapidly.

[73]See also Figure E1.

2025 years. It also appears that being treated in the Mid cohort decreased citations relative to being treated Late. And, finally, being treated Early versus Late resulted in a positive effect for a few years in the 21st century, although most years the sign of the effect is not identified.

### 5.2 The Effects of Court Mentions on Originalism

In our second application, we estimate the effect of judicial mentions on the reception for the originalist thinkers mentioned above. Recall that we separated the originalist works into four treatment cohorts: the Never-Treated cohort received zero court mentions; works in the Early treated cohort first received court mention in 1974–2009; works in the Mid treated cohort first received court mention in 2009–2019; and works in the Late treated cohort first received court mention in 2020–2025.

Our first originalist comparison is the Mid treated cohort versus the Never-Treated cohort. As shown in Figure 10, the upper and lower bounds on the $ATT$ for Mid treated units are shown in dashed-yellow and solid-green, respectively. Because originalism is a much younger legal movement than legal realism, the number of years each work has existed is much smaller. Perhaps unsurprisingly, this results in the sign of the $ATT$ being identified for only one year and only small values of the tuning parameter $C$. Notice, moreover, that the magnitude of possible $ATT$ values is much smaller: the $y$-axis range is much smaller than in the realist results above.

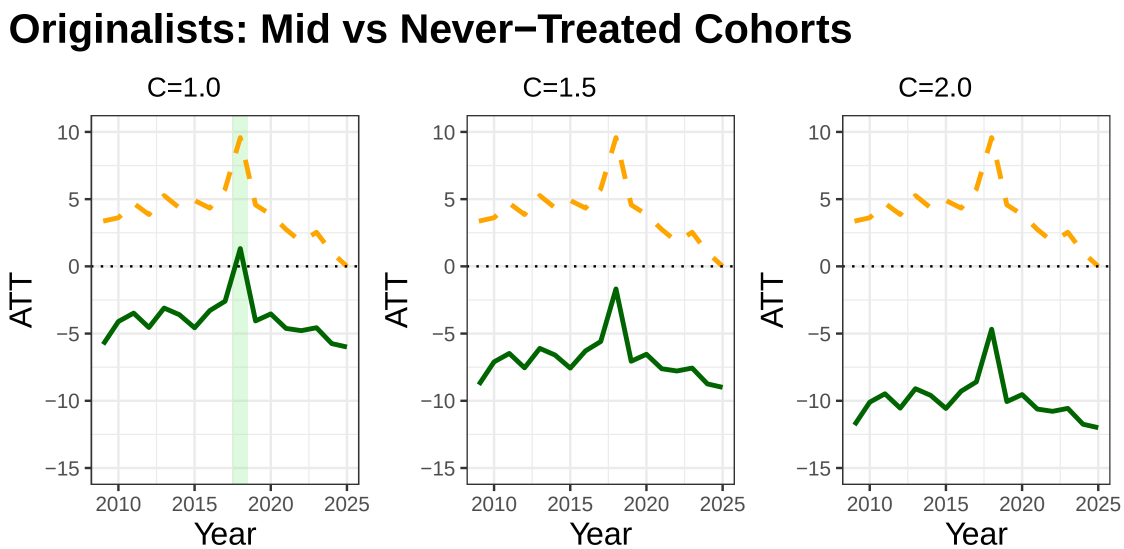


**Figure 10.** Realists: Mid vs Never-Treated Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

The comparison between Late treated originalist works and Never-Treated originalist works is similar. In Figure 11, we do not identify the sign of the effect from being in the Late treated cohort rather than the Never-Treated cohort. This means we cannot say whether judicial mention increased or decreased the academic citations for Late treated originalist works relative to never being mentioned in court opinions. We do show, however, that the magnitude of the effect is small either way: for $C = 1.0$ in the left-most panel, for example, the average effect of a Late judicial mention is less than 5 in absolute terms. This means, for example, that the data is consistent with judicial mention increasing (or decreasing) the academic citations to Late treated originalist works by some amount less than 5.

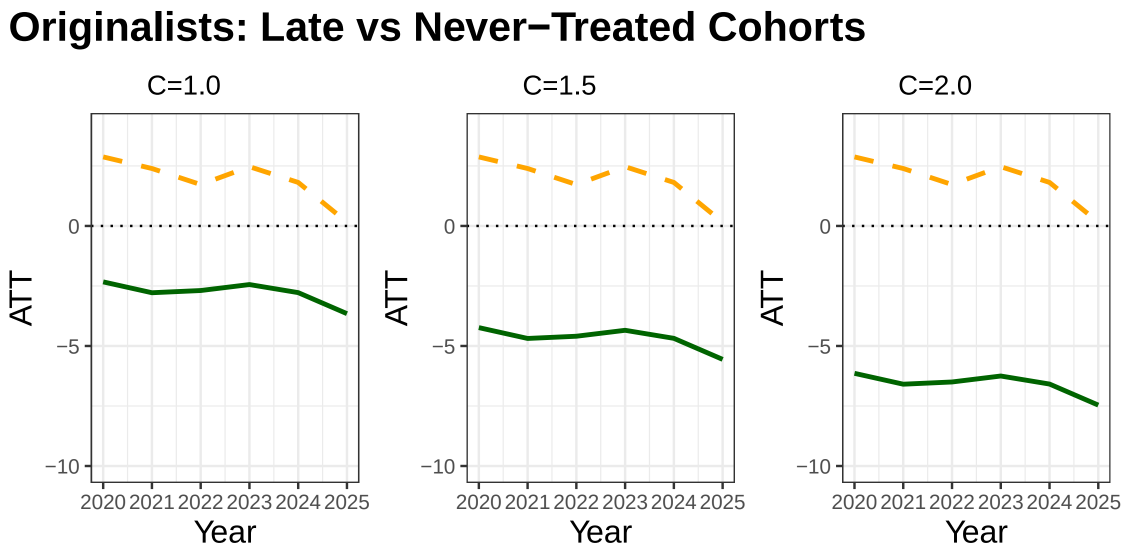


**Figure 11.** Originalists: Late vs Never-Treated Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

Again, we emphasize that we are working with data for only a small number of years.[74] Given the realist results, it appears that the impact of judicial mentions on academic citations takes many years (sometimes many decades) to noticeably affect academic citations. This makes sense given the time it takes for academics to appreciate particular works, process judicial changes to the law, and determine the canonical or important works in a particular legal tradition. In this case especially, originalist changes to constitutional law are active and ongoing. The effects of judicial originalism on legal academics may therefore take more time to appear in the data.[75]

### 5.3 Individual Treatment Effects

Our causal strategy can be easily adapted to estimate individual (unit-level) treatment effects.[76] We can consider, for example, comparisons between two works alone. In other words, we can relax the simplifying assumptions about the path space to bound

[74] Particularly for the Late vs Never-Treated cohorts.

[75] As in the realist case, see Appendix E for an additional result plot. There we compare Mid and Late treated originalist works and similarly find no sign-identification for $C$-values of $1.0, 1.5, 2.0$.

[76] As in the synthetic control literature. See, e.g., Abadie et al. (2010) and Abadie et al. (2015).

unit-level causal effects.[77]

In this case, one can bound individual-level treatment effects using just two works. Consider the example in Figure 12. In the left panel, the observed JSTOR citations for Holmes' "The Path of the Law" and Gray's *The Nature and Sources of the Law* are plotted blue and red, respectively. For 38 years, neither work receives any judicial mentions.

After 1947, however, the treatment paths for the two works diverge, which is reprsented by the black dotted line. After that time, the two works receive very different amounts of court mentions. The levels of the court mentions are shown by the volume of the blue and red dots for each work. As shown in the left panel of Figure 12, Holmes' "The Path of the Law" receives many more judicial mentions (larger blue dots). Using our Bounded Deviations approach for individual units, we can bound the causal effects from Holmes' work experiencing a different treatment path than Gray's work. The blue-striped region in the left panel is the counterfactual region, as discussed above.

[77]Recall as discussed above that few of the legal works we study experience the exact same treatment paths (judicial mentions) over the entire time. Instead, the different works are mentioned at different times and at different intensities (levels). For example, some legal works are mentioned early by many judges but then receive few (if any) judicial mentions in the late-20th or early-21st century. By contrast, some works only begin to receive a few judicial mentions many years after initial publication. And a select few works (e.g. *Corbin on Contracts*) have received high-levels of judicial mentions most years since publication. Our approach can allow for the potential popularity for each legal work to differ depending on when and how often judges mention the work in (published) court opinions. See further discussion in Appendix B.

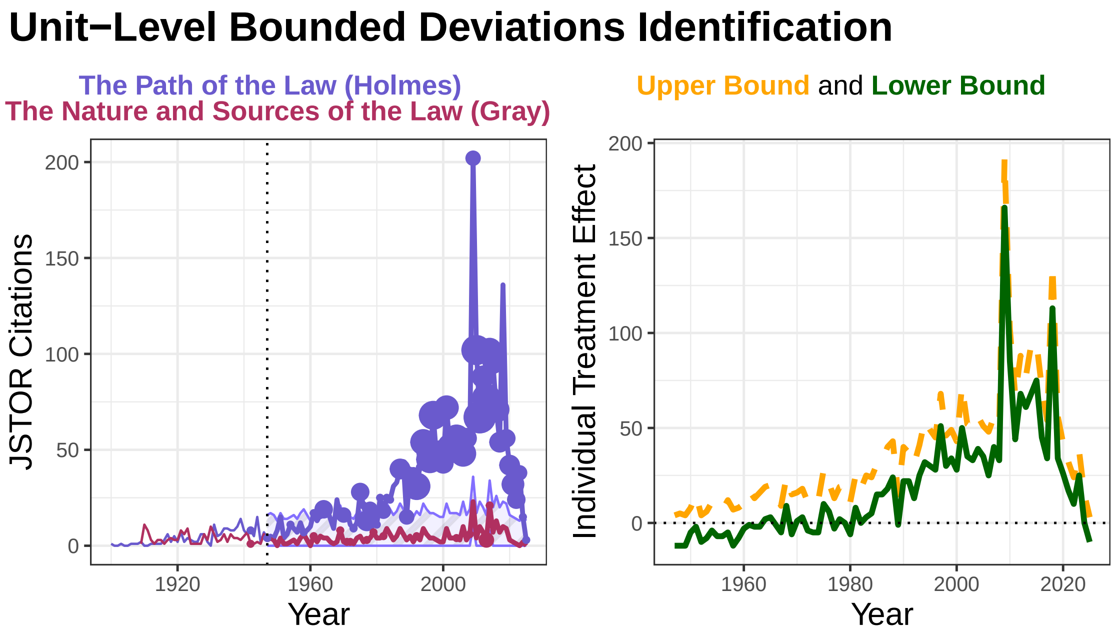


**Figure 12.** Individual (Unit-Level) Treatment Effects: Example

*Note:* The left panel shows the observed academic citations for two works. The dots represent years in which the work received court mentions, with the volume of the dots corresponding to more mentions. The right panel shows the individual treatment effect upper and lower bounds for $C = 1$.

In the right panel of Figure 12, we use the counterfactual region to compute the upper and lower bound on the individual treatment effect. As you can see, the sign of the treatment effect in earlier years includes the zero effect possibility. However, in later years the treatment effect is positive. This is because the citations to Holmes' work are much higher than the top part of the counterfactual region. See Appendix B for further discussion.

## 6 Conclusion

In our application, we show how courts impact the popularity of legal theories. Specifically, we show how a bounding approach can estimate the impact of judicial mention on citations to realist and originalist works. Although we find it is probably too soon to say with certainty how judges affect originalist citations, we find positive effects of judicial mentions on citations to treated legal realist works relative to works that are never mentioned by judges.

Causal inference in this setting is important because existing empirical scholarship has primarily focused on how judges merely *use* academic legal works, but has not estimated the *causal effects* of judges on legal academia. Another contribution of our work is, therefore, to bring tools in the causal inference revolution to legal questions. We show how to estimate counterfactuals for particular works if they had experienced alternative treatment paths. Importantly, we flexibly allow for heterogeneity across treatment paths and units, while keeping the causal identification problem tractable.

Estimating counterfactuals in settings where potential outcomes may depend on the entire history or path of treatment requires these simplifying assumptions. As discussed above, the general strategy is to assume how unobserved counterfactual quantities evolve to determine the causal effects of treatment at different times (cohorts). The Bounded Deviations method we propose is a simple and flexible way to do this. It bounds counterfactual paths by using information from units with similar treatment histories until some divergence point. After the divergence point, we use the maximum deviation from the pre-divergence periods to bound the counterfactual path.

Our approach can be applied to further legal settings. As shown above, boundedness assumptions nest the standard parallel trends assumption and are robust when counterfactual trends are not parallel. This is important because much empirical legal work relies on the difference-in-difference strategy to infer causal effects. Violations of the assumptions underpinning that strategy can result in failure to isolate the sign or magnitude of the effect of treatment.

Finally, the data in this paper is useful for isolating descriptive trends. We show large variation in when (and to what degree) judges mention works in two important legal movements. Some works are mentioned early and often, whereas others are mentioned (if at all) long after publication. The mechanism by which certain works assist (or not) judges in deciding cases is of interest to legal academics interested in the development of particular legal subjects. Further studies may use the vast amount of text data in judicial opinions and law review articles to better understand what makes academic legal scholarship important, popular, and useful.

## References

Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association 105*(490), 493–505.

Abadie, A., A. Diamond, and J. Hainmueller (2015). Comparative politics and the synthetic control method. *American Journal of Political Science 59*(2), 495–510.

Baude, W. (2015). Is originalism our law? *Columbia Law Review 115*, 2349–2408.

Black, R. C. and J. F. Spriggs II (2013). The citation and depreciation of u.s. supreme court precedent. *Journal of Empirical Legal Studies*.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Chilton, A. S. and E. A. Posner (2015). An empirical study of political bias in legal scholarship. *Journal of Legal Studies 44*.

D'Amato, A. (1989). Can any legal theory constrain any judicial decision? *University of Miami Law Review 43*, 513.

D'Amato, A. (1999). The effect of legal theories on judicial decisions. *Chicago Kent Law Review 74*, 517–527.

Detweiler, B. T. (2020). May it please the court: A longitudinal study of judicial citation may it please the court: A longitudinal study of judicial citation to academic legal periodicals. *Legal Ref. Servs. Q. 39*.

Diamond, S. S. and P. Mueller (2010). Empirical legal scholarship in law reviews. *Annual Review of Law and Social Science 6*, 581–599.

Edwards, H. T. (1992). The growing disjunction between legal education and the legal profession. *Michigan Law Review 91*(1), 34–78.

Feldman, A. (2018). Empirical scotus: With a little help from academic scholarship. *SCOTUS Blog*.

Fitzpatrick, B. T. and P. K. Varghese (2017). Scalia in the casebooks.

Fronk, C. R. (2010). The cost of judicial citation: An empirical investigation of citation practices in the federal appellate courts. *Journal of Law, Technology, and Policy*.

Goldsmith-Pinkham, P. (2024). Tracking the credibility revolution across fields. *Arxiv*.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Greene, J. (2009a). On the origins of originalism. *Texas Law Review 88*(1).

Greene, J. (2009b). Selling originalism. *Georgetown Law Journal 97*.

Gulati, M. and V. Sanchez (2002). Giants in a world of pygmies? testing the superstar hypothesis with judicial opinions in casebooks. *Iowa Law Review 87*, 1141.

Kennedy, D. and W. W. Fisher (2006). *The Canon of American Legal Thought*. Princeton University Press.

Leiter, B. (2005). *The Blackwell Guide to the Philosophy of Law and Legal Theory*. Wiley.

Leiter, B. (2010). *American Legal Realism*. Wiley.

Manski, C. F. (1997). Monotone treatment response. *Econometrica 65*(6), 1311–1334.

Manski, C. F. (2007). *Identification for Prediction and Decision*, Chapter 9. Harvard University Press.

Manski, C. F. and J. V. Pepper (2018). How do right-to-carry laws affect crime rates? coping with ambiguity using bounded-variation assumptions. *The Review of Economics and Statistics 100*(2), 232–244.

Newton, B. (2012). Law review scholarship in the eyes of the twenty-first century law review scholarship in the eyes of the twenty-first century supreme court justices: An empirical analysis. *Drexel Law Review 4*, 399–416.

Posner, R. A. (1992). The deprofessionalization of legal teaching and scholarship. *Michigan Law Review 91*, 1921.

Posner, R. A. (2016). *Divergent Paths: The Academy and the Judiciary*. Harvard University Press.

Pound, R. (1908). Mechanical jurisprudence. *Columbia Law Review 8*(8), 605–623.

Rambachan, A. and J. Roth (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 2555–2591.

Rumble, W. E. (1965). The paradox of american legal realism. *Ethics 75*(3), 166–178.

Sawyer III, L. E. (2018). Principle and politics in the new history of originalism. *American Journal of Legal History 57*(2), 198–222.

Sirico Jr., L. J. and B. A. Drew (1991). The citing of law reviews by the united states courts of the citing of law reviews by the united states courts of appeals: An empirical analysis. *University of Miami Law Review 45*, 1051–1060.

Tamanaha, B. Z. (2008). Understanding legal realism. *Texas Law Review 87*, 731.

Thalken, R., E. H. Stiglitz, D. Mimno, and M. Wilkens (2023). Modeling legal reasoning: Lm annotation at the edge of human agreement. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing,*, 9252–9265.

Whittington, K. (2004). The new originalism. *Georgetown Journal of Law and Public Policy 2*(2), 599–613.

# APPENDICES

# A Appendix A: Additional Descriptive Plots

Appendix A includes additional descriptive plots. We include the following plots: two additional academic citation series for the legal realist works (cumulative vs yearly citations); court mentions for realist works yearly; and court mentions for each author aggregated across works and for each decade.

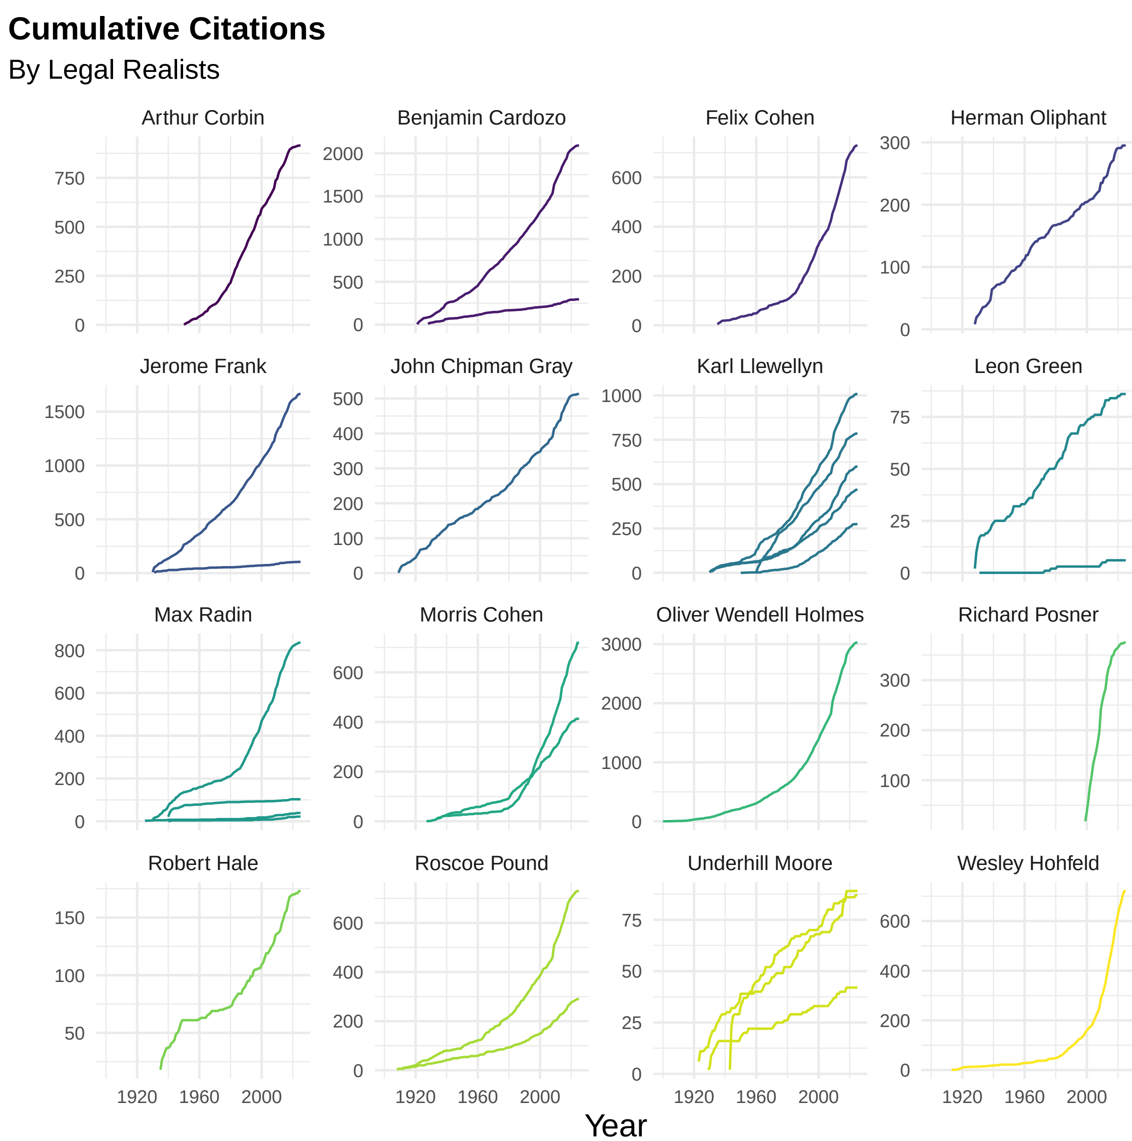


**Appendix Figure A1.** Cumulative Citations

*Note:* The $y$-axis differs for each legal realist. The different lines for each legal realist are for their different texts.

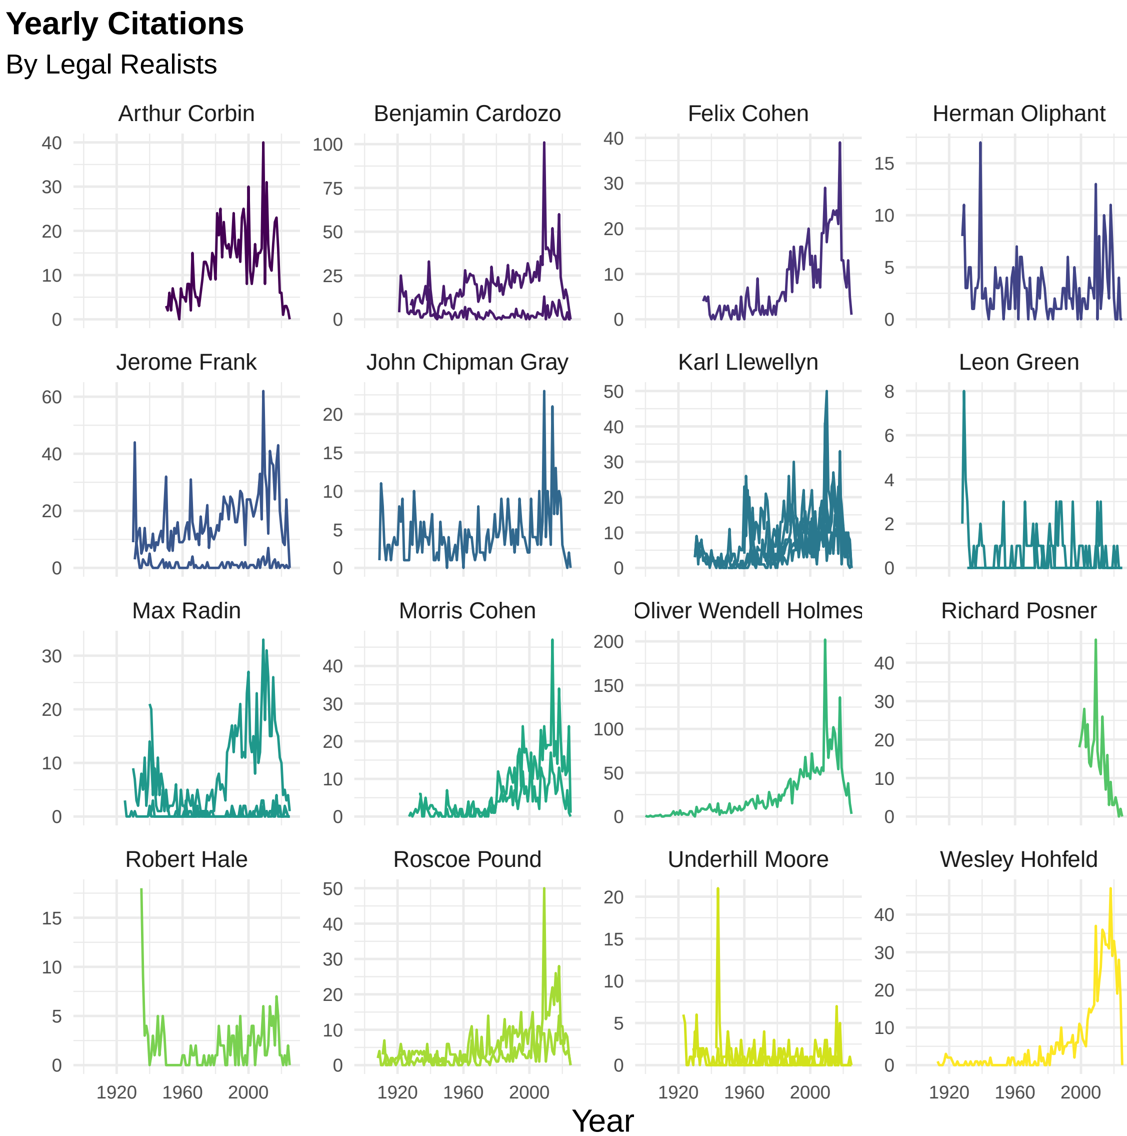


**Appendix Figure A2.** Yearly Citations

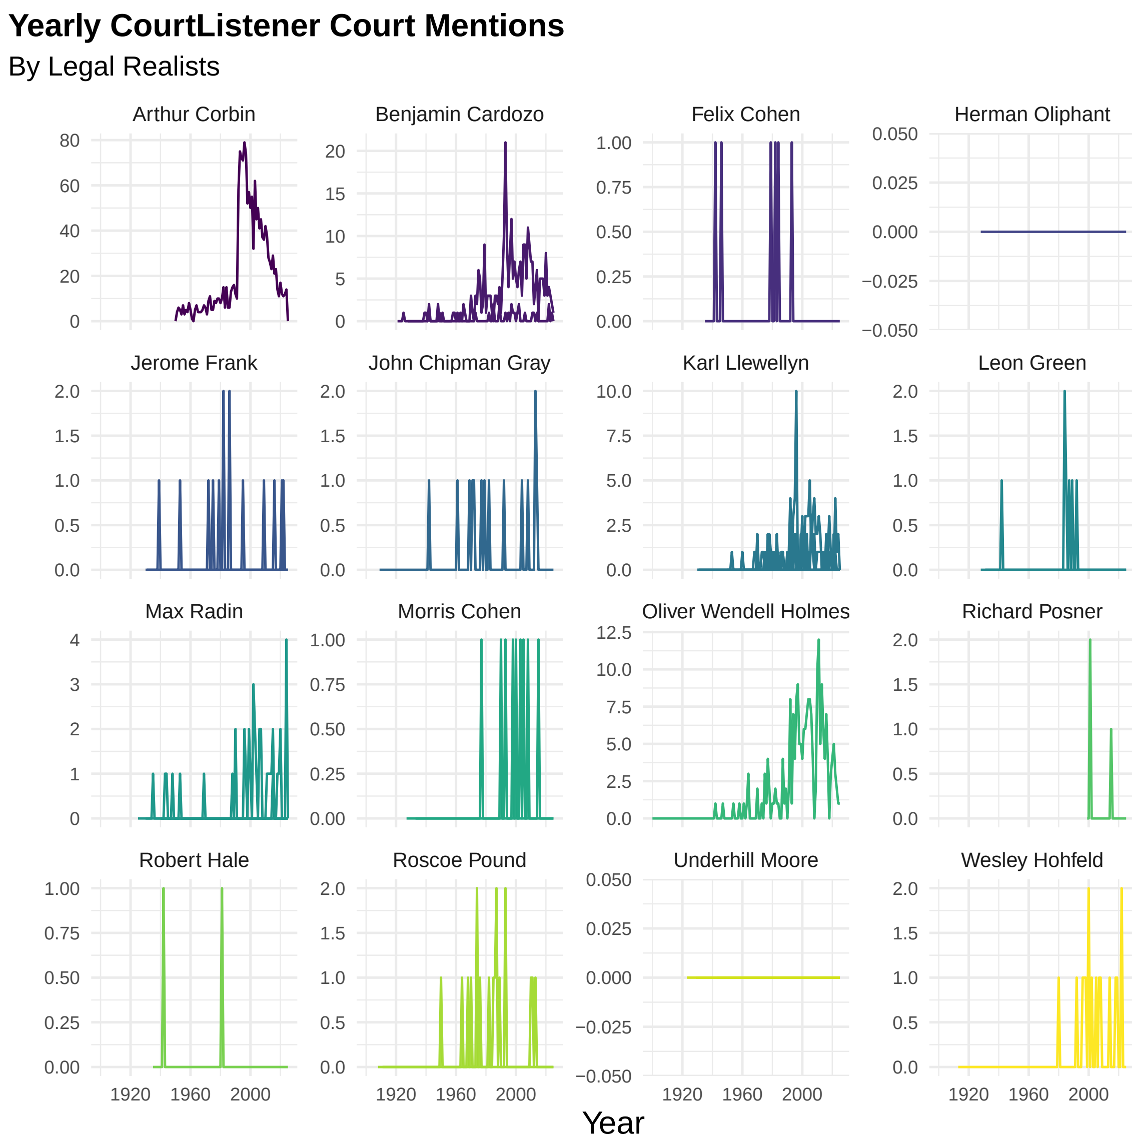


**Appendix Figure A3.** Yearly CourtListener Court Mentions

*Note:* The $y$-axis differs for each legal realist.

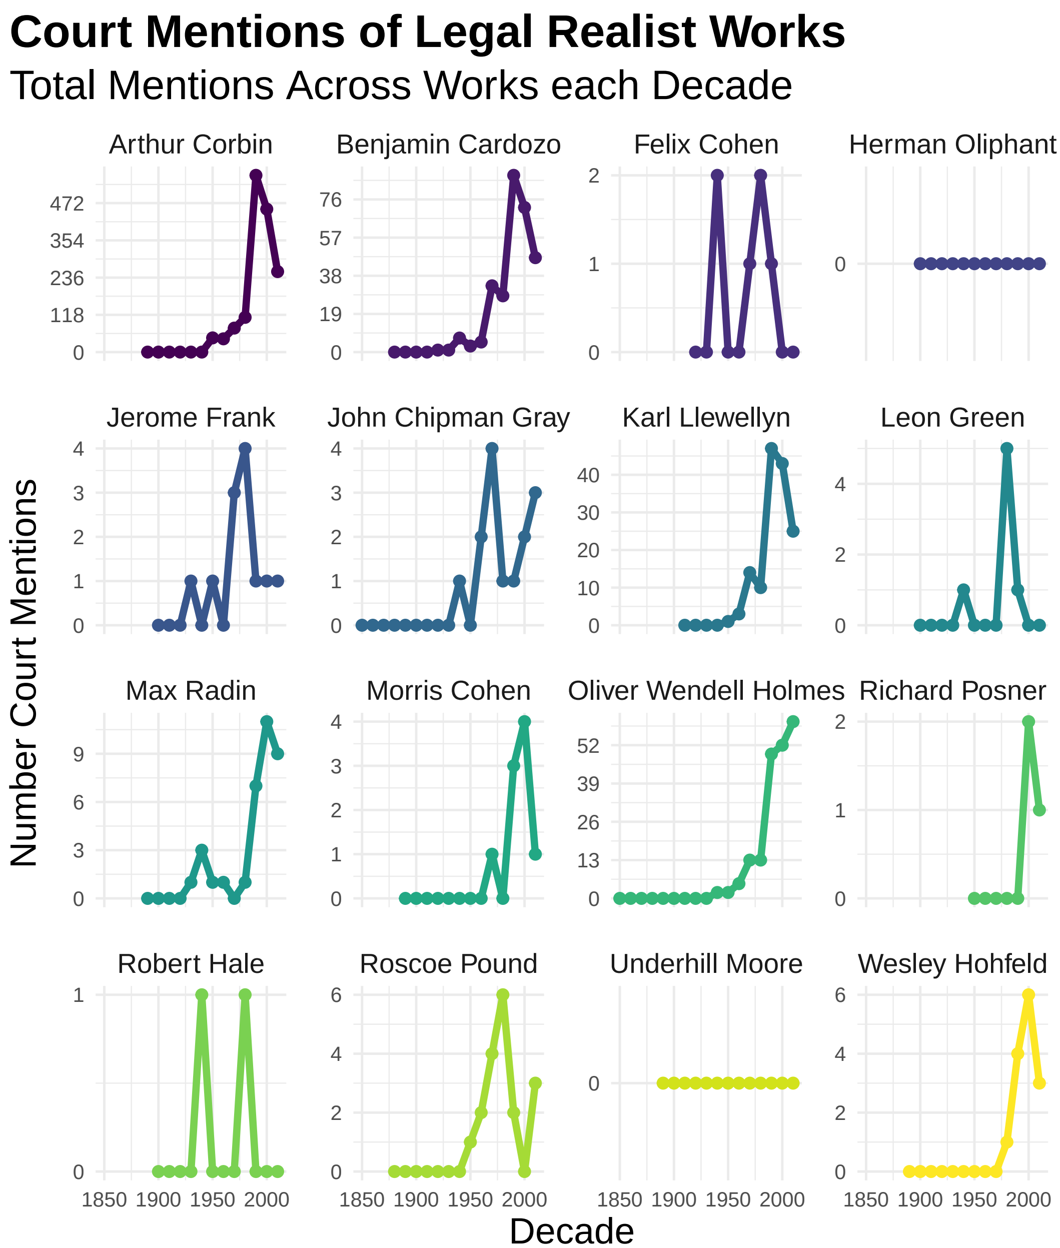


**Appendix Figure A4.** Legal Realist Court Mentions

# B Appendix B: Counterfactual Complexity & The Curse of Dimensionality

Appendix B includes discussion and a plot of how potential outcomes grow as a function of the complexity of treatment paths. As mentioned in-text, researchers typically rely on simplifying assumptions rather than allowing for full path heterogeneity. Here, we discuss two types of assumptions that are (implicitly or explicitly) used in practice: memory and anticipation assumptions.

The first set of assumptions we term memory assumptions, which restrict how the potential outcome in period $t$ can depend on treatments received prior to time $t$. Formally, we define *memory of degree* $m$ to be the assumption that, for all $i$ and in each time $t$, we have:

$$Y_{it}(\vec{k}) = Y_{it}(k_{t-m}, \ldots, k_t, \ldots k_T)$$

This restricts the potential outcomes in period $t$ to depend only on the treatment path up to $m$ periods prior to $t$. For example, memory of degree zero (*no memory*) means that potential outcomes in time $t$ depend only on treatments received during or after time $t$.

The second set of assumptions are anticipation assumptions. Commonly used the DID literature, this assumption restricts how potential outcomes in time $t$ can depend on treatments received after time $t$. Formally, we define *anticipation of degree* $a$ to be the assumption that, for all $i$ and in each time $t$, we have:

$$Y_{it}(\vec{k}) = Y_{it}(k_1, \ldots, k_t, \ldots k_{t+a})$$

This restricts the potential outcomes in period $t$ to depend only on the treatment path up to $a$ periods after $t$. For example, anticipation of degree zero (*no anticipation*) means that potential outcomes in time $t$ depend only on treatments received during or before time $t$.

Figure B1 illustrates the dimensionality problem. We plot the number of potential outcomes under different memory assumptions and no anticipation. We consider settings with ten time periods and between 1-10 treatment dosages other than the baseline "no treatment" dosage. As seen in Figure B1, stronger memory assumptions greatly reduce the dimensionality of the problem. However, as the cardinality of the treatment set $\mathcal{K}$ grows, the total number of potential outcomes remains large: under no memory and no anticipation, there are 100 total potential outcomes for each unit.

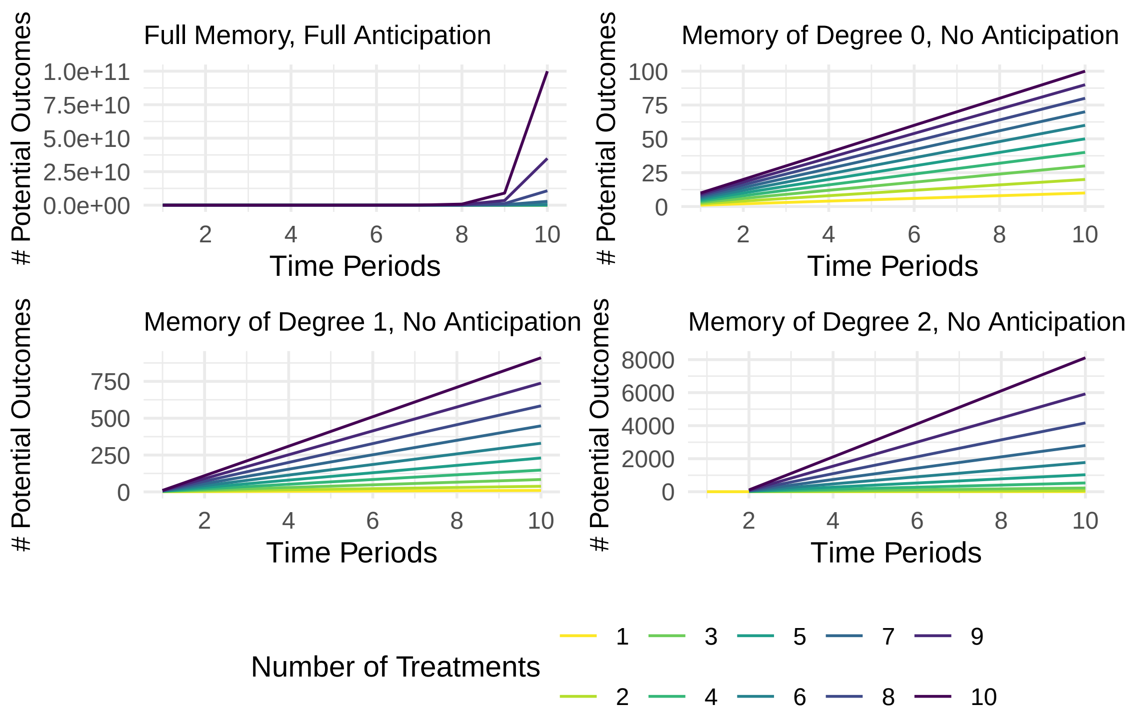


**Appendix Figure B1.** Potential Outcomes: Combinatorial Explosion

*Note:* This plot shows the total number of potential outcomes as a function of time periods ($x$-axis) and the number of treatments (line color). The number of treatments is the number of treatments in addition to a baseline "no treatment" dosage. The yellow line (1 treatment) is, e.g., a binary treatment in each period.

Even with these restrictions on the nature of the treatment path, counterfactual imputation is not as straightforward as the simpler setting. For one, there are now more counterfactual paths. Researchers might ask what the unit's outcome would be if it were not treated at all, if it were treated earlier, or if it were treated later. The standard approach—comparing treated units to not-yet-treated units to obtain treatment effects—mimics the simpler settings described above.

These complications in dynamic settings are exacerbated as the complexity of the treatment path increases. Imputing counterfactuals with complex treatment paths is challenging for two reasons. The first is sparsity in the path space itself. Indeed, as shown in Figure 2, there is limited overlap in the treatment histories for units that are mentioned at least once in judicial opinions. The second is that even if many paths are observed, there will be too few observed units on each path to feasibly point identify a causal estimate. Both these make the usual machinery employed in causal inference difficult to work in these settings. Under such complex paths, there is no clear demarcation between treatment and control units, particularly if all units get some treatment at some point.

Because it works to impute unit-level counterfactuals, the Bounded Deviations method we use is one possible solution. Our identification strategy partially identifies counterfactual paths for units receiving different treatment paths. If the length of a treatment path grows, the number of such counterfactual paths can be unmanageable quickly. To make this manageable, we could only limit to paths that we observe, or make parametric assumptions. For example, we could define a distance metric on treatment paths and assume that counterfactual outcomes are a function of the distance between paths. Counterfactual paths that are closer to the observed treatment path, e.g., may only have small, predictable deviations from the observed outcomes.

### B.1 Possible Problems with Dimension-Reduction Assumptions

Memory-anticipation assumptions clearly reduce the dimensionality of the identification problem. Still, computational and interpretational problems may remain, especially in settings with numerous treatment levels (i.e. doses). But the main problem is that such assumptions can introduce bias and are difficult to motivate or justify in settings without well-founded theory. That is, if there are meaningful differences in potential outcomes along different treatment paths, comparing units along different paths can bias causal estimates.

Consider, for example, the simple two-period model with binary treatments available both time periods. If there is genuine anticipation, then the counterfactual averages along paths $(0,0)$ and $(0,1)$, say, may differ. That is, units that are untreated at time $t=1$ but treated in time $t=2$ may anticipate receiving treatment, act on it, and thereby experience different outcomes than units that are never treated (i.e. whose treatment paths are $(0,0)$). If we then assume there is no anticipation, i.e. that $Y(0,0) = Y(0,1) \equiv Y(0)$, we improperly "mix" these units together. This can result in estimates of $\mathbb{E}[Y(0)]$ that differ substantially from the true population counterfactual averages $\mathbb{E}[Y(0,0)]$ and $\mathbb{E}[Y(0,1)]$. The direction and magnitude of the bias from assuming no anticipation can be difficult to predict or correct *ex post*.[78]

In our setting, modeling anticipation may be especially difficult. Although researchers studying, e.g., the effects of particular pre-announced, well-discussed congressional or executive legal actions may be able to model anticipation, it is not clear

[78] This problem is analogous to the problem of "forbidden comparisons" in staggered adoption settings. See, e.g., Goodman-Bacon (2021); see also Callaway and Sant'Anna (2021).

what legal academics anticipate in terms of judicial citations. It may be very difficult to anticipate which academic articles judges will cite, yet some law review pieces are prepared with particular court cases (esp. Supreme Court cases) in mind. Indeed, some legal academics may expect their writings to be successful because of pending judicial decisions and even file *amicus* briefs encouraging judges to engage with their scholarship. Particularly for a legal movement like the originalism movement in constitutional law, judges seem to be openly inviting academic discussion of particular legal issues.[79]

Simplifying the treatment path space is, therefore, not without possible downsides. However, in spaces with high-dimensional (time-varying) treatments, simplifying from full-path treatment heterogeneity to a few treatment cohorts is necessary to obtain overlap. As stated above, the flexible partial identification assumptions we propose above can work to impute unit-level counterfactuals (as in, e.g., synthetic control). This is because the unobserved potential outcomes depend on the observable potential outcomes but differences are "bounded."

### B.2 Individual Treatment Effects & Identifying Assumptions

Finally, we briefly sketch how one could consider individual treatment effect estimands. As in the synthetic control literature,[80] we consider individual treatment effect estimands of the following form:

$$\Delta_i(t, \vec{k}, \vec{k}') = Y_{it}(\vec{k}) - Y_{it}(\vec{k}')$$

In what follows, we can obtain similar bounds on these individual treatment effects by the modified assumption below:

**Assumption 2** (Individual Bounded Deviation). *Let $T_{div} = argmin_t(k_t \neq k'_t)$ for two units $i$ and $j$ on distinct treatment paths $\vec{k}$ and $\vec{k}'$, respectively. The following holds for each $t \geq T_{div}$ and some $C \in \mathbb{R}_+$:*

$$\left| Y_{it}(\vec{k}') - Y_{jt}(\vec{k}') \right| \leq C \cdot \max_{t < T_{div}} |Y_{it} - Y_{jt}|$$

[79] See, e.g., *Franklin v. New York*, 604 U.S. (2025) (Alito, J., concurring) (discussing the possibility of overturning Confrontation Clause precedent based on "[h]istorical research").

[80] See, e.g., Abadie et al. (2010) and Abadie et al. (2015).

# C Appendix C: Originalist Treatment Histories

Appendix C contains the treatment histories for each of the originalist works studied here. For each work, we plot the number of judicial mentions the work receives each year following publication. The levels are colored according to the legend at the top of each figure.

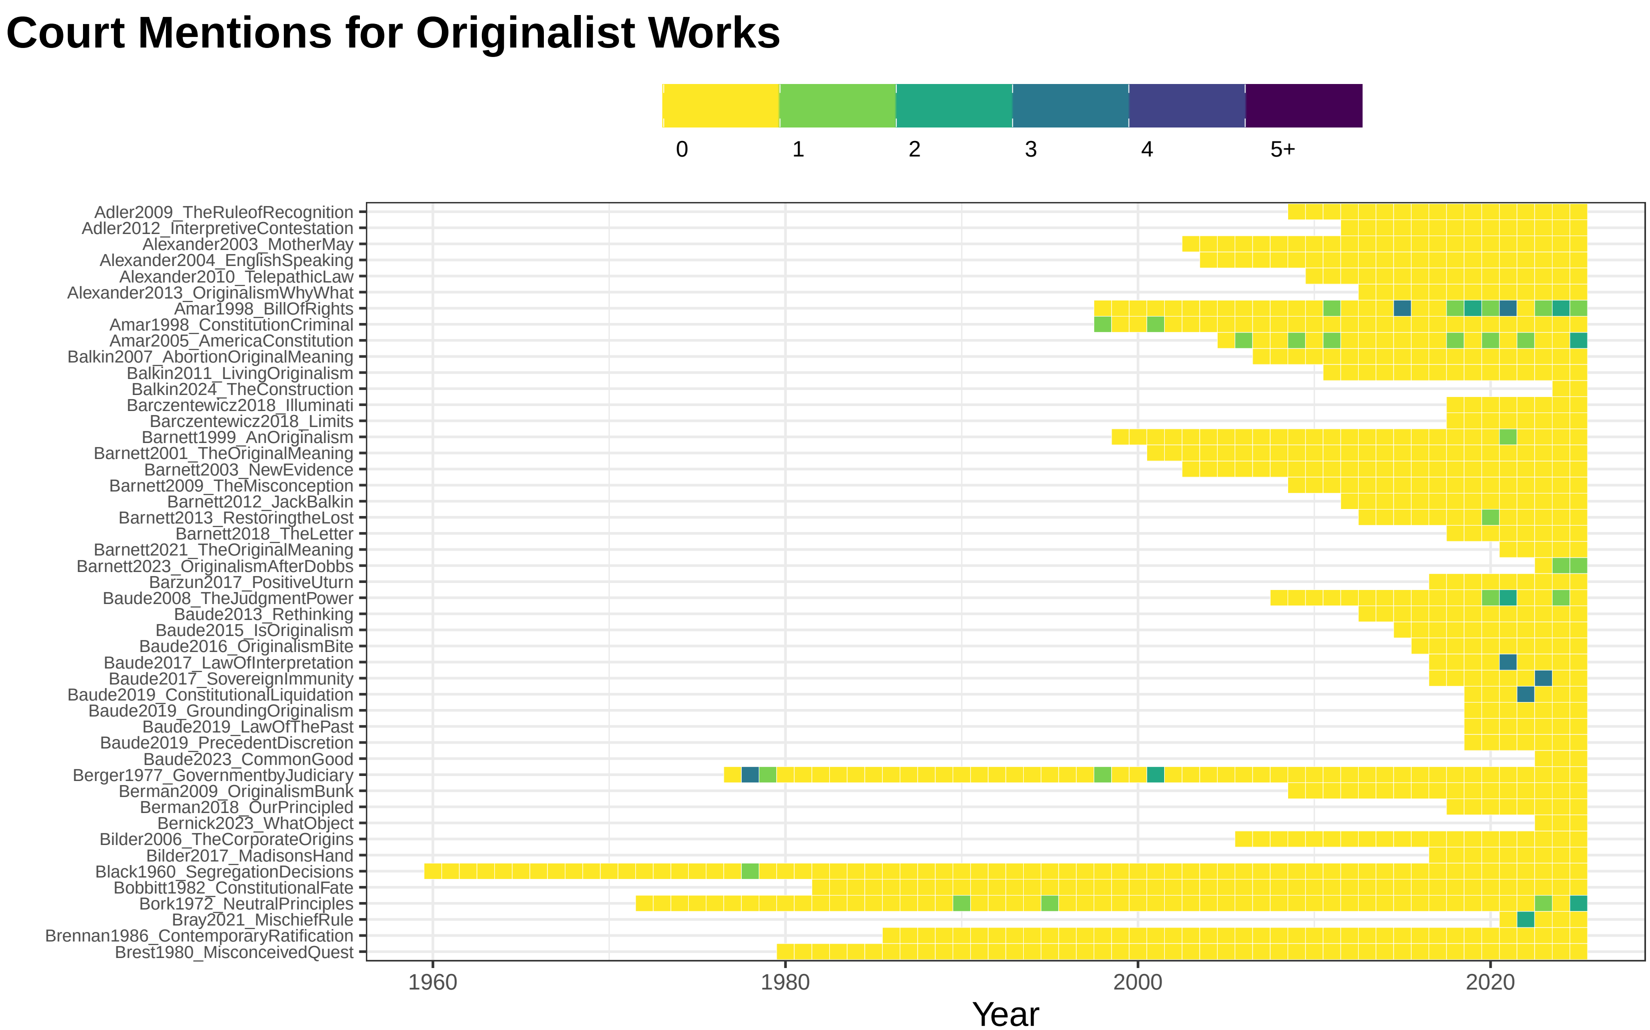


**Appendix Figure C1.** Originalists: Court Mentions (Part 1)

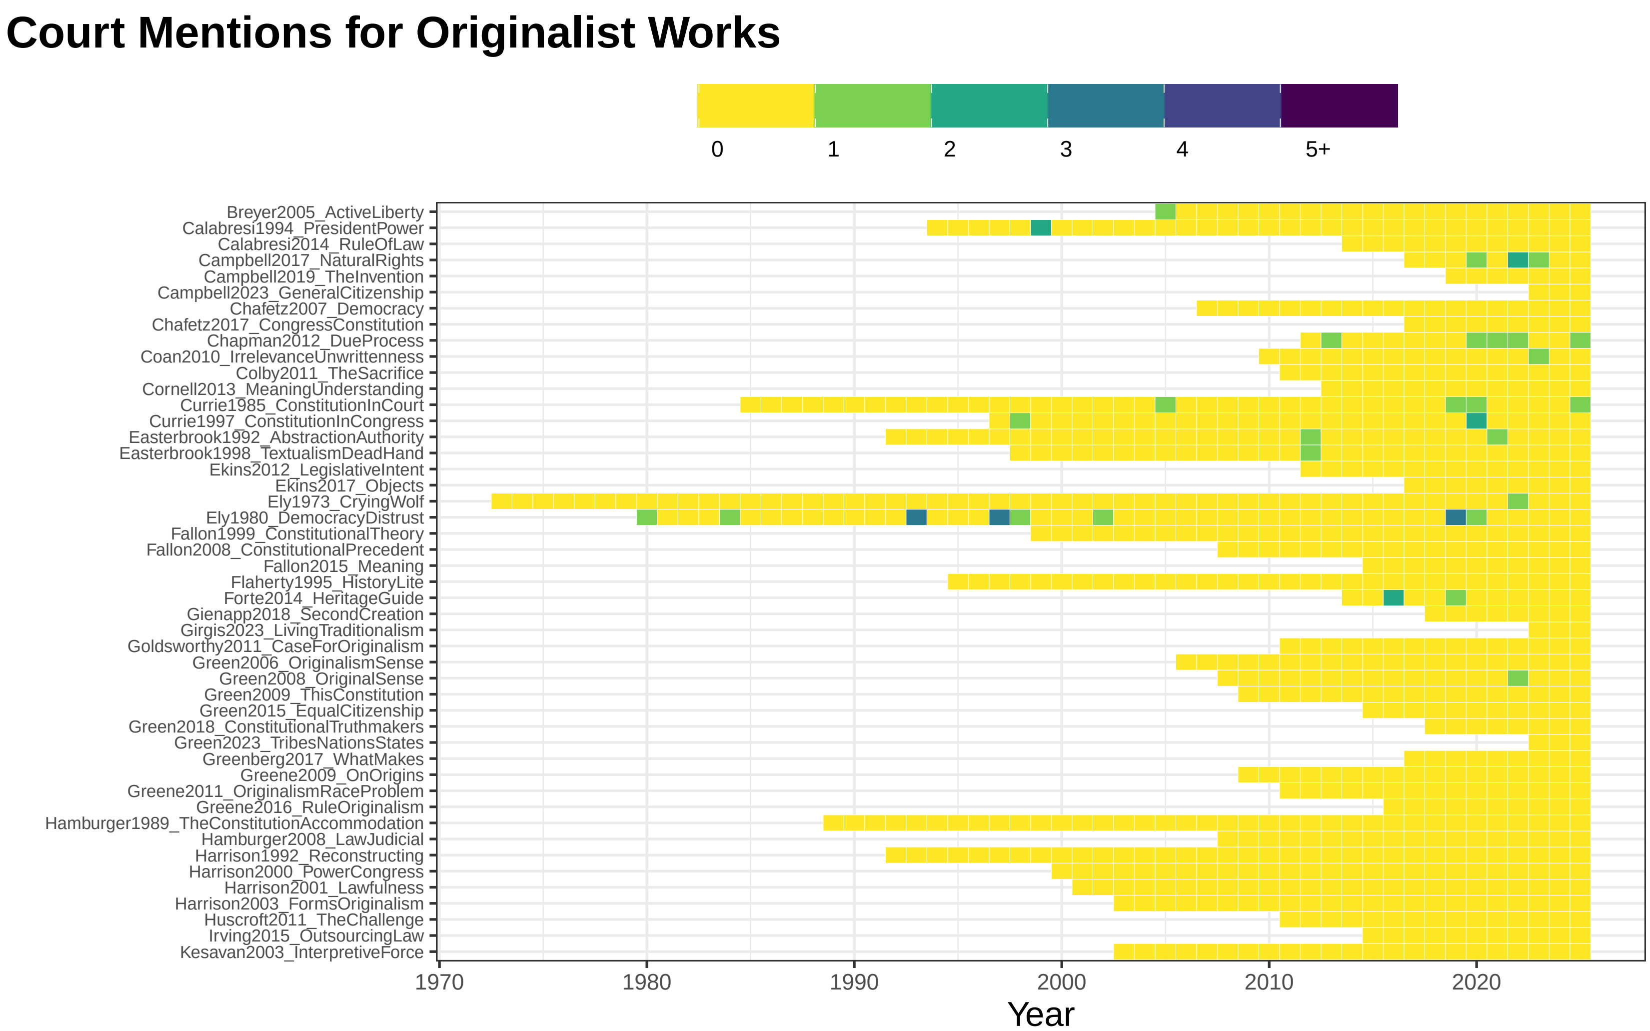


**Appendix Figure C2.** Originalists: Court Mentions (Part 2)

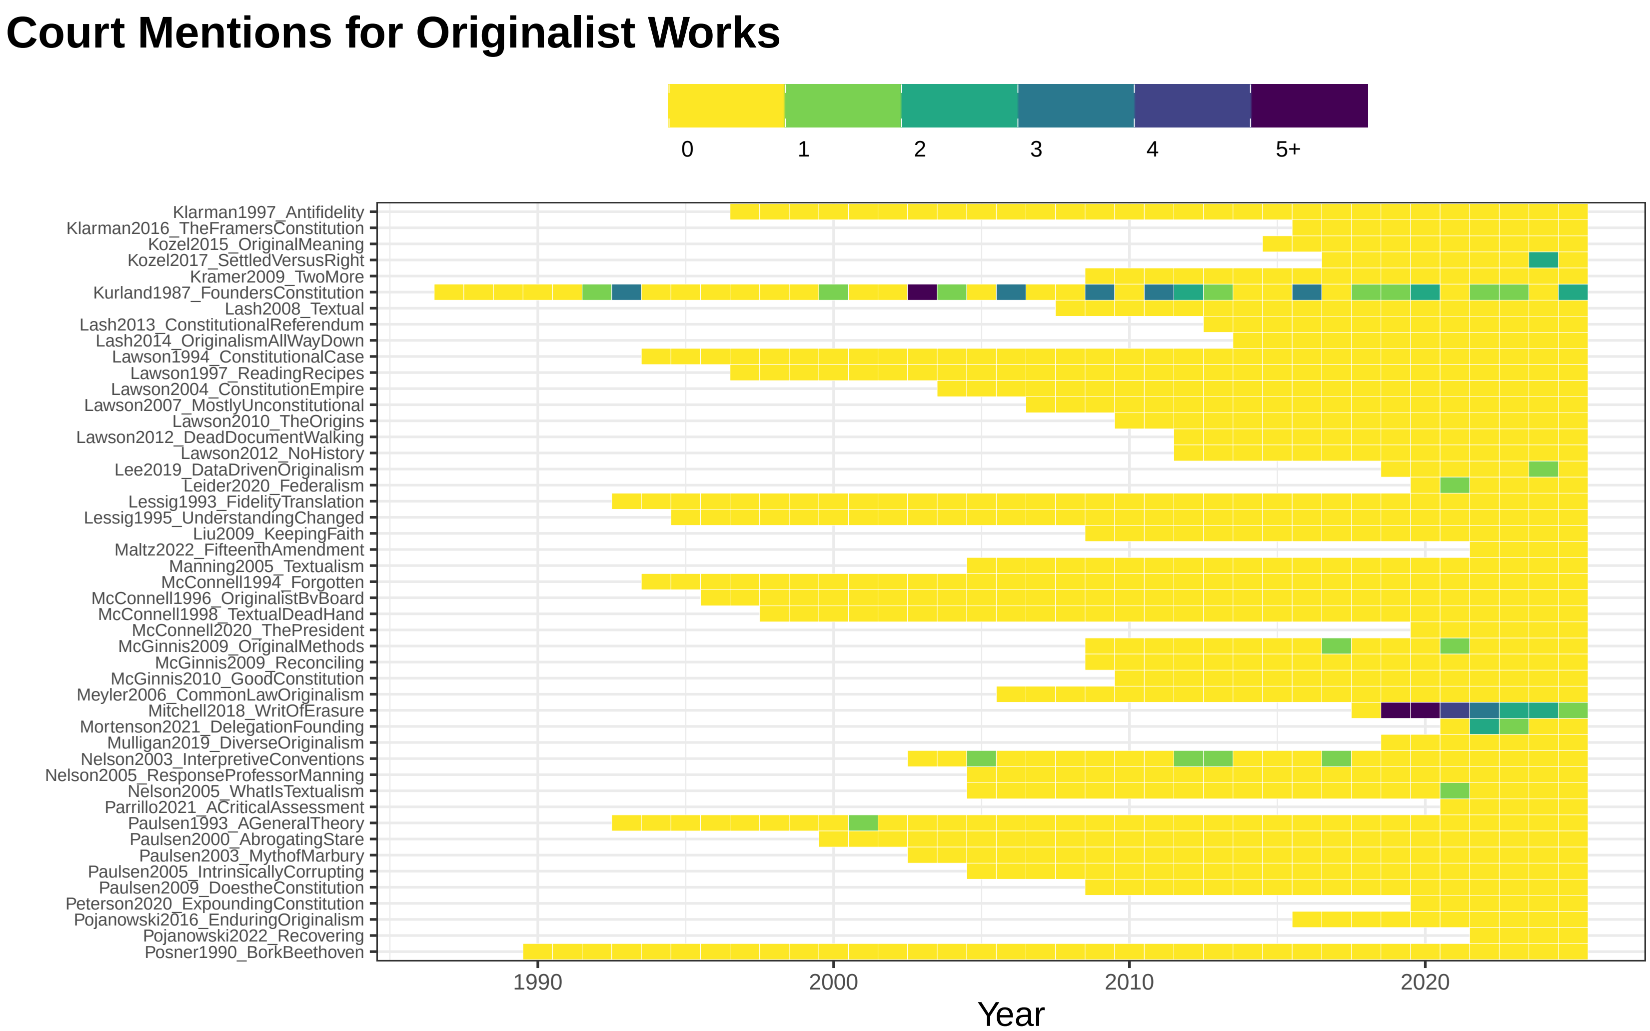


**Appendix Figure C3.** Originalists: Court Mentions (Part 3)

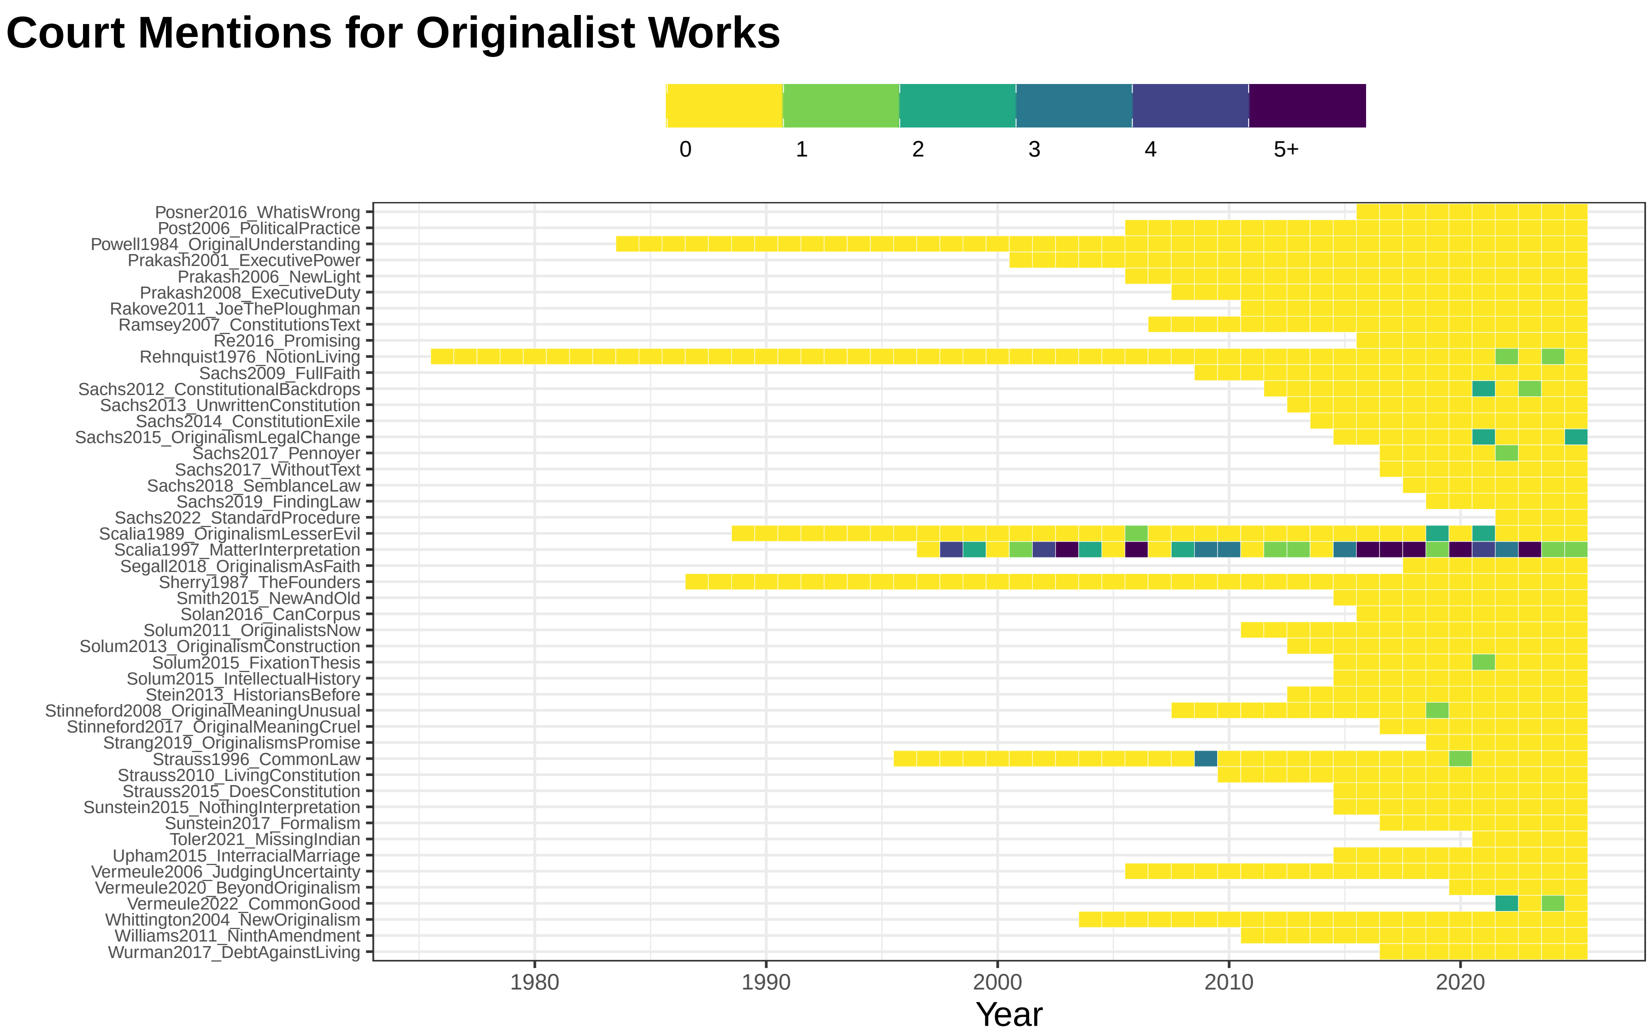


**Appendix Figure C4.** Originalists: Court Mentions (Part 4)

# D Appendix D: Bounded Deviations Tuning Parameter

Appendix D show how the Bounded Deviations identification strategy depends on the tuning parameter $C$. We show how this identification tuning parameter affects the upper and lower bounds in Figure D1. The upper and lower bounds are plotted in the right and left panels, resp., for different values of $C$. In the plot, we consider varying $C$ from being very small (yellow colored lines) to very large (dark blue lines). Smaller $C$ values result in the upper and lower bounds being very close. For example, the yellow lower bound and yellow upper bound lines do not differ by much. By contrast, large $C$ values result in larger counterfactual regions. This is shown by how far apart the lower and upper bounds are when $C = 3$ (i.e. dark blue lower bound and dark blue upper bound are far apart).

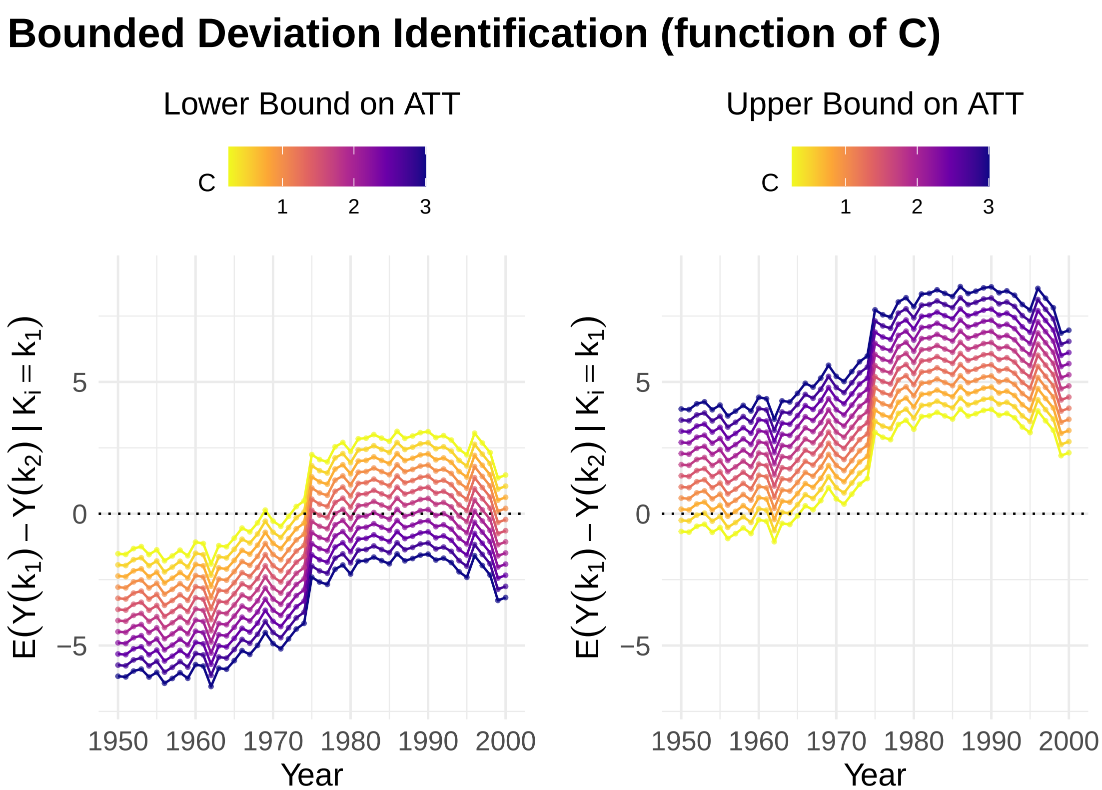


**Appendix Figure D1.** ATT Bounds as a Function of $C$

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$.

As shown in Manski and Pepper (2018), researchers can study how sensitive their causal estimates are to this tuning parameter.

# E Appendix E: Additional Results Plots

Appendix E contains additional results plots. First, we show the cohort average JSTOR citations for the realists and originalists. For both groups, the never-treated averages are, in most periods, lower than the three treatment cohort averages.

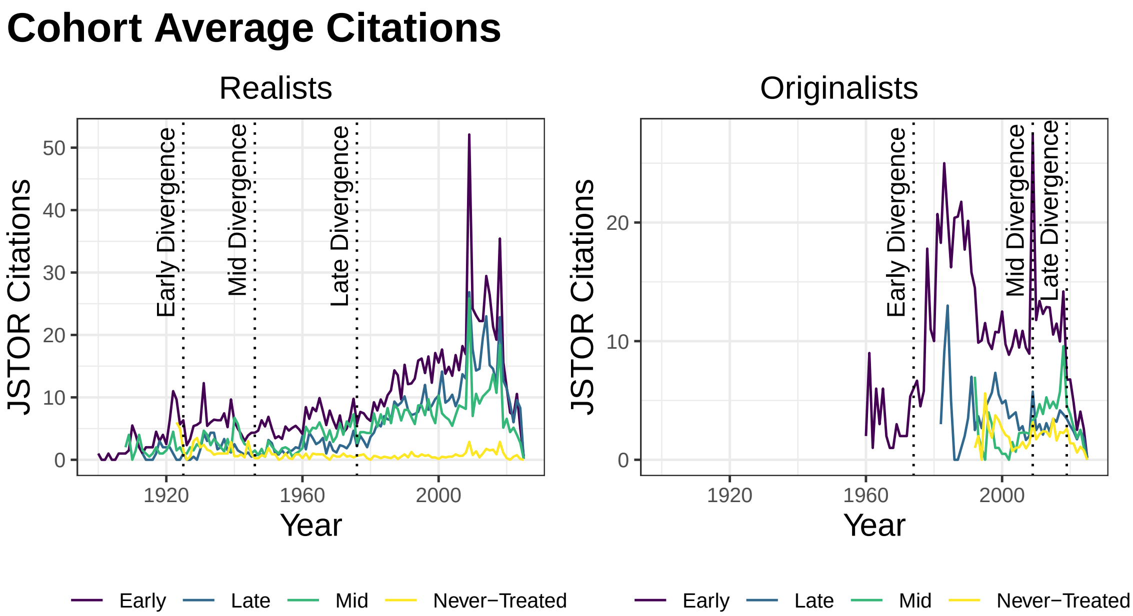


**Appendix Figure E1**

*Note:* The divergence periods that demarcate the treatment cohorts are plotted as vertical black dotted lines.

The remaining figures are ATT bounds plots as above for different values of the tuning parameter $C$. Please note in each plot the cohorts that are being compared. Unlike the $ATT$ bounds discussed in-text, these are comparisons between treated cohorts.

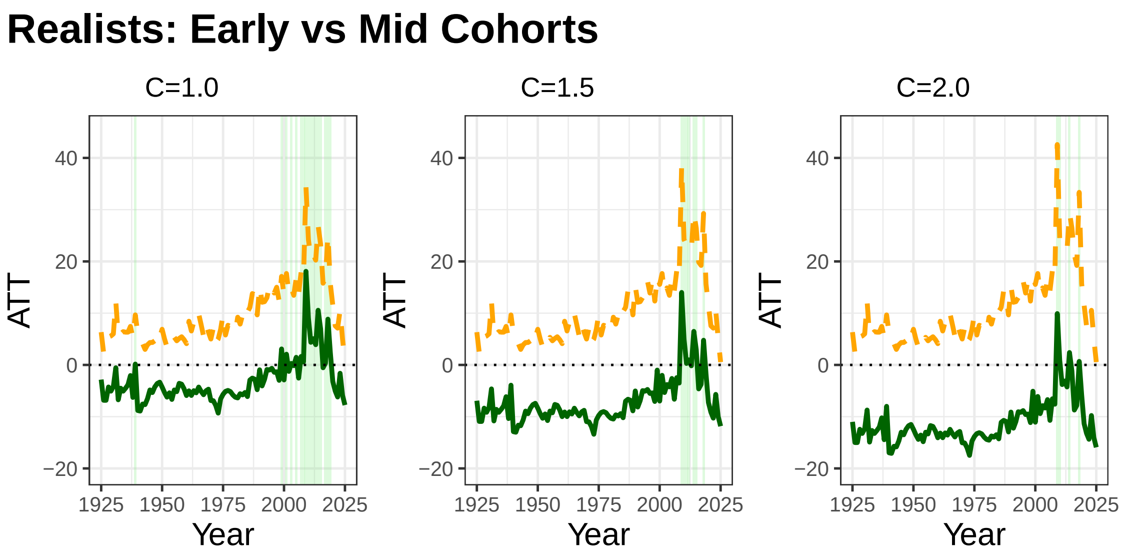


**Appendix Figure E2.** Realists: Early vs Mid Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

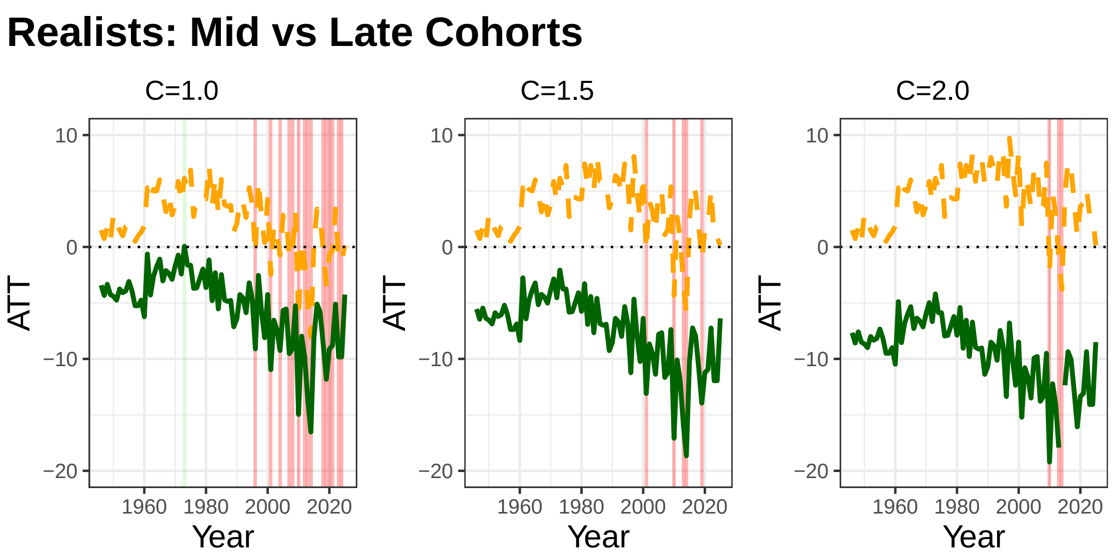


**Appendix Figure E3.** Realists: Mid vs Late Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

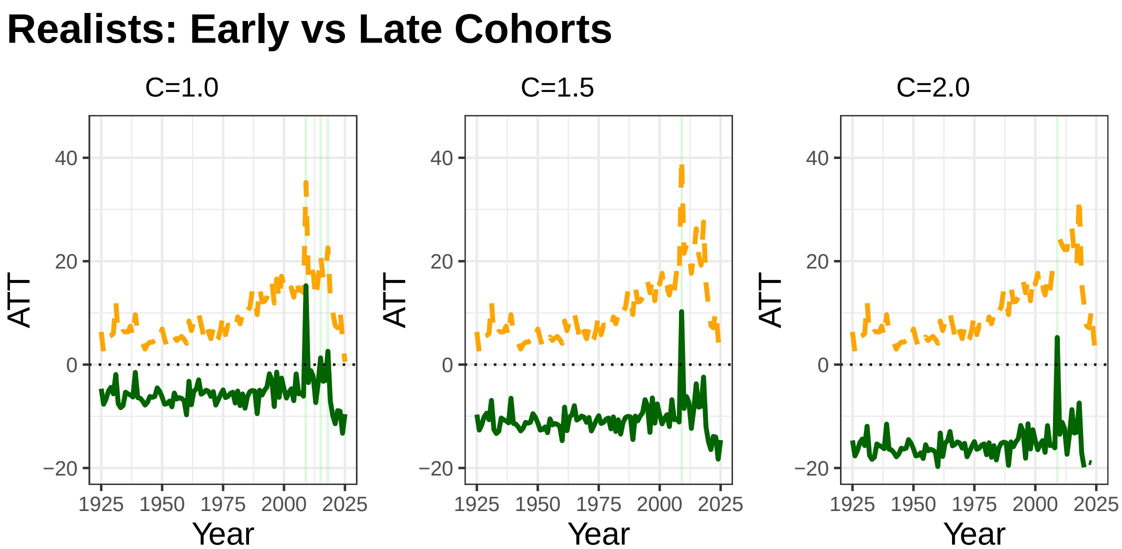


**Appendix Figure E4.** Realists: Early vs Late Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

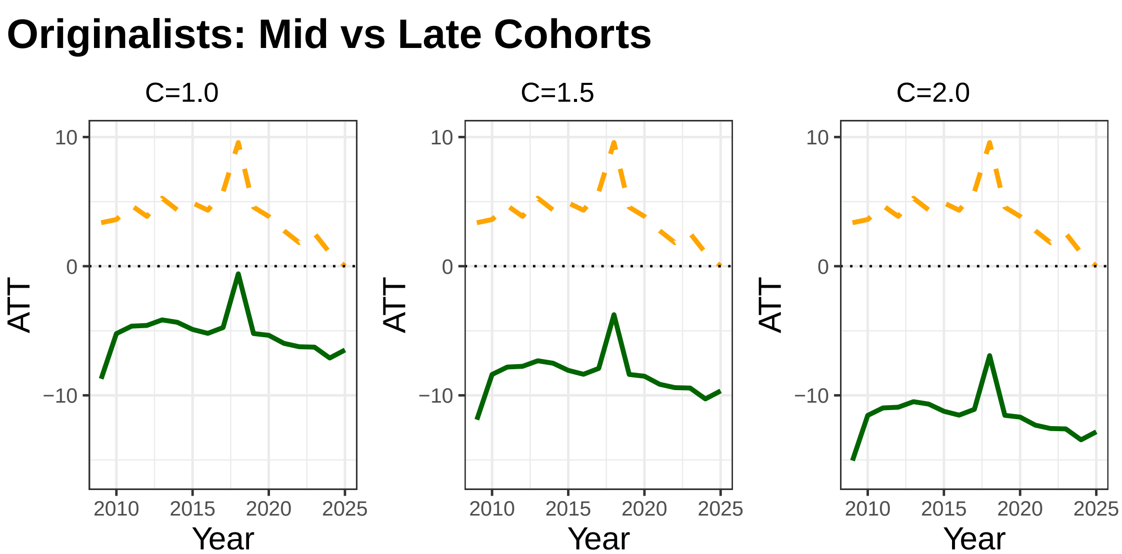


**Appendix Figure E5.** Originalists: Mid vs Late Cohorts

*Note:* In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the $ATT$ upper bound is the dashed-yellow line, and the $ATT$ lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.