Measuring Institutional Innovation: Recombinant Novelty in Early-Modern English Caselaw

Peter Grajzl† Peter Murrell‡

August 20, 2024

Abstract

Using only the words in texts, we measure institutional innovation by integrating topic-modeling estimates with an algorithm that identifies innovative pairings of ideas. Our measure captures a well-accepted conception of innovation—recombination. We apply our method to early-modern English legal case-reports, producing a first-ever, yearly time-series of institutional innovation for that paradigmatic context. We generate both aggregate timelines, and timelines disaggregated by type of court and area of law. The generated data offer the promise of a canonical, quantitative record of legal development, facilitating a re-examination of many aspects of conventional wisdom on English history. We revisit the timing of institutional innovation and the role of the common law. We re-assess the legacy of major historical figures that have been regarded as major contributors to caselaw development. We contrast our measure of innovation with that based on case citations. Our method could be applied to any text corpus.

Keywords: novelty, recombination, caselaw, early-modern England, Edward Coke, citations

JEL Classifications: K40, N01, N44, O35, B40, C81, D02, P48

*This paper was completed while Murrell was Hallsworth visiting professor at the University of Manchester. We are grateful to participants at seminars at Manchester and Oxford for helpful comments. We thank Dan Johnson, Kathryn Hendley, and Norbert Hornstein for valuable advice, and Rishi Paladugu, Awaid Yasin, and Yicheng Zhang for research assistance, the latter two under the auspices of the Policy Analytics Research Program at the Mercatus Center.
† Department of Economics, The Williams School of Commerce, Economics, and Politics, Washington and Lee University, Lexington, VA 24450, USA and CESifo, Munich, Germany. Email: grajzlp@wlu.edu
‡ Department of Economics, University of Maryland, College Park, MD 20742, USA. Email: pmurrell@umd.edu

> Combinatory play seems to be the essential feature in productive thought.
> Albert Einstein, 1945
> Let us now peruse our ancient authors for out of the old fields must come the new corn. Edward Coke ca. 1632

1. Introduction

The power of the recombination of existing elements in producing novelty is widely recognized, from the adaptive advantage of sexual reproduction, to scientific progress, to new constitutions, and to the dynamics of firms. In the social sciences, and especially in history, the conception of novelty as recombination has provided limited empirical leverage because of difficulties in measurement. The small number of existing studies usually use data with a pre-existing classification of observations. This data feature readily appears when studying patents and academic publications, but not in other settings, thereby currently confining large-scale empirical studies of recombinant novelty to a limited domain. We address this problem by integrating topic-model estimates into an algorithm that identifies innovative pairings of ideas. This provides a general method for producing time-series of recombinant innovation from texts alone.

As we argue in Section 2, one broader area where measuring novelty via recombination might provide fundamental insights is in the history of ideas. An abiding question of history is when the world began the process of modern development (see, for example, North et al. 2009, Ogilvie and Carus 2014, Murrell 2017, Henriques and Palma 2023). To the extent that ideas can both drive and reflect change, one approach to identifying transformative historical periods would entail finding when novel recombinations of existing ideas occurred. We apply this insight to study legal change in early-modern England, a time and a place where democracy, the rule of law, and modern economic growth first began to appear together.

One important constructive characteristic that set England apart was its legal system, especially its judge-made caselaw. During the early-modern period, advances in English law occurred primarily within caselaw rather than via statute. Our objective then is to find those time periods when surprising, novel recombinations of previously separate legal ideas appeared most frequently in caselaw. We thereby produce a host of new facts on legal development: when it occurred; in which areas of law and courts it was most prominent; which judges and law reporters contributed most to novelty; and how novelty relates to the citation record

The primary challenge in accomplishing our aims is that analysis must begin with raw text data since the legal-historical record appears in case-reports. These reports intermingle many legal ideas, but these legal ideas come unlabeled and are not classified by area of law. Yet, identifying instances of recombination requires observing the melding of ideas that fall under separate and easily-discernible classifications. This is an insight that has driven empirical studies on novelty using non-law-related data from current times by using, for example, referenced journals in science articles or technology classifications in patents (Uzzi et al. 2013, Wang et al. 2017). To address this empirical challenge in our legal-historical context, we draw on the estimates generated by

Grajzl and Murrell (2021a), who applied topic modeling to the core corpus of early-modern legal case texts, the English Reports (Renton 1900-1932). Grajzl and Murrell's (2021a) estimates apportion case-reports among endogenously identified topics, or classifications. Although all topics appear in all reports to some extent, we create a method of classifying reports by focusing on only the most prominent topics in each report. Section 3 below provides the necessary details on Grajzl and Murrell (2021a), as well as explaining how we identified the prominent topics within each report.

Section 4 then uses the method of Wang et al. (2017) to assign a recombinant novelty score for each report. The essence of our application of their algorithm is to identify case-reports that feature a pairing of prominent topics that has not appeared in reports from previous decades. Such a pairing is a novel recombination of existing elements. We assign higher novelty scores to reports that contain a greater number of novel pairings. But the algorithm also discounts recombinations that were less surprising or easier to produce, as assessed by whether the two topics had been combined with similar sets of topics in the recent past.

Since we generate an entirely new measure of early-modern English legal development, reflecting only the words that reporters wrote, we provide validation, in Section 5. We show broad similarity between the summary statistics for our data and those from studies of novelty using modern non-legal corpora. We examine a sample of cases and find consistency between the details of our novelty assessment and the content of case-reports. We show that our estimates of legal development are broadly similar to those present in timelines of the references in Blackstone's Commentaries (Blackstone 1893 [1765]), Holdsworth's (1938 [1903-1938]) history, and in Baker's (2019) textbook. In addition, we note that temporal variations in aggregate novelty scores can be easily understood in light of broadly accepted ideas about English political and legal history.

Our central contributions are three-fold. First, we show how to quantify a well-accepted notion of innovation that relies on categorizations of texts, even though the legal case-reports are uncategorized. Our method would be eminently applicable to any corpus whose documents span a long time period. Second, we generate both aggregate and disaggregated time-series data on innovation in English caselaw, providing entirely new insights into the historical record. Third, we show how our produced data can be used to examine many aspects of the conventional wisdom on English history, including the role of the common law, the pace and timing of early-modern English legal development, and the legacy of prominent historical figures who have contributed to the creation of caselaw. Thus, Sections 6-10 provide a host of new facts on legal development in early-modern England. Section 11 brings together our major findings.

2. Empirical recombination and case law

That recombination of existing ideas is a prime ingredient of innovation is widely accepted. This notion appears broadly in the history of science, for example in the works of Koestler, Kuhn, Popper, Feyerabend, and many scholars of more recent vintages. Testaments from successful scientists on their own methods abound, as suggested by the epigraph of this paper (Hadamard 1945). Schumpeter (1934) views recombination as central in entrepreneurship. Nelson and Winter

(1985: 130) characterize recombination as the essence of innovation in firms, and indeed in all creativity, in art, science, and practical life. Support for such views is widespread and can be found in analyses of scientific publications, law, patents, music, and video game production.[1]

Recombination is a natural fit for the study of innovation in caselaw (Whalen et al. 2017). Generally, conceptual recombination is triggered when distinct concepts, doctrines, and rules come into play simultaneously to address a problem that needs to be solved (Thagard 1988). Consider a puzzling legal case: events on the ground lead to an unusual confluence of circumstances, necessitating a decision that requires the combining of previously separate, well-established aspects of law, leading to an innovative ruling or novel reasoning. The new combination of circumstances on the ground is then exactly analogous to the Morley-Michelson experiment on the speed of light that provoked the relativity revolution in physics. This view fits well with the philosophy of the common law as practiced in early modern England. Judges "liked to think of their law as an unchanging body of common sense and reasoning…If human reason does not change, the law cannot change; it is only the application of old ideas to new situations which creates the appearance of change" Baker (2019: 206). Indeed, judges sometimes conjured up the new situations for themselves, by the device of legal fictions, which allowed the importing of law from a different domain than was directly implied by the facts of a case. Thus, via the recombination of older ideas judges could be innovative without imparting the appearance of developing wholly new doctrine.

Legal 'change' in the early modern period provides illustrations. Assumpsit, an action to recover damages from contractual breach, was in the 16th century widely applied in claims arising from wrongs that resulted in physical damage. Around the same time, money claims appeared in actions on debts that had been explicitly acknowledged. Thus, the question arose as to what would happen if the debt were implicit, as in contract breach. The resolution involved a new rule that led to the reading into contracts of a promise to pay the sum due, recombining elements already in assumpsit, actions on debt, and contract law (Lieberman 1989: 127-8). Nearly a century later, promissory notes obtained the legal characteristics of money, through the combination of caselaw on merchant custom, assumpsit, negotiability of bills of exchange, and conveyance of title (Fox 1996). Then in the 18th century, the strict settlement was created to control the futures of family assets. This combined previous legal ideas on trusts and on inheritance to create a system that governed the dispositions of the great landed estates for two centuries (Harris 2000: 149). In sum, recombination lies at the heart of legal innovation.

3. Data

We draw on the topic-model estimates generated by Grajzl and Murrell (2021a), where a full description of the pertinent model, methods, and estimation can be found.[2] The use of topic models rather than alternative machine-learning methods is appropriate in the current context because the

[1] In addition to references below, see also de Vann et al. (2015), McMahan and McFarland (2021), and Mueller (2002).
[2] A less technical description, using the same topic-model, appears in Grajzl and Murrell (2022). The dataset is available at http://www.econweb.umd.edu/~murrell/Data/ER/ER.html.

novelty algorithm leverages clearly identifiable subcategories of caselaw. It is the identification of these subcategories, that is, the topics, that lies at the heart of topic modeling. Grajzl and Murrell's source data are pre-1765 case-reports included in the definitive version of the English Reports (Renton 1900-1932; hereafter ER). After pre-processing, a 100-topic structural topic model is estimated (Roberts et al. 2016). This summarizes 52,949 case-reports in a 52,949-by-100 matrix. Element ij of the matrix indicates the proportion of case-report i that arises from topic j. Each topic can be construed as referring to a distinct set of ideas (e.g., Assumpsit; Pleadings on Debt; Intestacy). The use of such a document-topic matrix lies at the heart of all substantive applications of topic-modeling.

We adopt the topic names chosen by Grajzl and Murrell (2021a; Table 1), capitalizing to distinguish them. Many of these names would be familiar to legal historians, such as Precedent, Nuisance, Habeas Corpus, or Writs of Error. Others are more obviously a product of topic-model discovery rather than following traditional legal categories. Examples are Disentangling Heirs, Rulings on the Calendar, and Determining Damages & Costs. It is easy to see how such topics can arise out of the synthesis of many legal rulings, but one will not find them in standard legal classifications.

We focus on a subset of the Grajzl and Murrell (2021a) data, reducing the dimensions of their document-topic matrix in two ways. First, the number of words in case-reports varies considerably. Some case-reports consist of just a few sentences. In an initial implementation of our methodology applied to all 52,949 cases, we read a sampling of cases in conjunction with estimates of novelty scores. False positives for novelty appeared most commonly in short case-reports. In contrast, all inspected longer case-reports with high novelty scores had characteristics that indicated novelty. We thus focus our analysis on the subset of case-reports that had more than 500 words. In addition, we restrict our analysis to cases adjudicated after the year 1500, when reports begin to be more numerous.

Second, as DiMaggio et al. (2013) argue, one of the valuable features of topic modeling is that non-substantive words are consigned to residual topics, thus improving the quality of the substantive topics. For example, in the Grajzl and Murrell (2021a) dataset there is a topic named Non-Translated Latin. This topic was excluded from the analysis because its presence or absence does not indicate genuine legal novelty. Also, there are six topics reflecting various case-reporting styles. These were excluded from the topics used to assess legal novelty because changes in reporting style, while indicating novelty in form, do not indicate novelty in substance. As a result of these reductions in dimension, our analysis begins with a 14,523-by-93 document-topic matrix. Appendix A provides more details on the construction of our dataset. Figure A.1 shows the yearly distribution of case-reports.

4. Measuring novelty

Broad quantitative, empirical analyses of novelty focusing specifically on recombinant processes are rare compared to the many theoretical and descriptive studies that promote the idea of recombinant novelty without proceeding to measurement. For broad coverage of an era,

empirical analyses of recombinant novelty in any domain must access large comprehensive databases, which in no small part explains the focus in the existing literature on academic publishing and patents. And given that recombinant analyses produce demands on computer power that grow exponentially with dataset size, comprehensive studies have only become possible in recent years. Uzzi et al. (2013), a seminal study in this vein, characterizes journal articles from 1950 to 2000 using pairs of references in bibliographies. More novel articles are those that use reference-pairs that have been less used elsewhere. Wang et al. (2017) follow the broad ideas in Uzzi et al. (2013), while using a less computationally burdensome procedure. They examine all scientific publications in 2001 and characterize novelty using pairs of cited journals. Youn et al. (2015), Verhoeven et al. (2016), and Veugelers (2022) follow similar approaches using patent classifications. Kaplan and Vakili (2014) use topic modeling to identify breakthrough patents, defined as those for which a specific topic rises to a threshold level of prominence. Our method of measuring recombinant novelty integrates elements of all these approaches.

Finding the recombinant novelty of a case-report relies on identifying topic pairs that are prominent in that report and not prominent in previous ones. However, a core assumption of topic modeling is that all topics are present in all documents. Nevertheless, the prevalence of any given topic varies widely between case-reports. Therefore, we adopt a simple criterion for when a topic is prominent in a case report: when the topic proportion is two standard deviations above the mean topic proportion, with standard deviation and mean calculated separately for each topic proportion across all case-reports.[3] When a topic proportion meets this criterion, we refer to the topic as being 'in' the case-report, or, equivalently, prominent in the case-report. All classification methods applied to texts will involve an analogous procedure: for example, a patent classification refers to a prominent element in the patent description even though the underlying science will generally be drawn from many facets of existing knowledge.

For each case-report, we find all prominent topics and the set of pairs of these topics (e.g., four prominent topics gives rise to six distinct pairs).[4] For each of these topic pairs, we find whether it is 'new', that is, has not appeared in any case-report in the previous 30 years. A simple case-report-level measure of novelty could then be the count of new topic pairs: the number of new recombinations.

However, some novel topic pairs arise more naturally than others, indicating relatively less novelty. A topic pair is especially innovative if the pairing could not have been easily foreseen on the basis of past topic combinations. Thus, for each novel topic pair, we discount novelty based on how easy it was to create the novel topic pair. Specifically, suppose that in a given case-report a new topic pair consists of topics k and l. We find two vectors v_k and v_l, each of dimension 91 (the total number of topics, 93, minus these two), with a typical element v_{pq} equal to the number of

[3] Readers familiar with text processing will recognize the analogy to the notion of term frequency–inverse document frequency applied to a document-vocabulary matrix.
[4] Hence, our procedure now follows the approach of Wang et al. (2107) except that we use topic pairs when they use journal citation pairs.

case-reports in the last 30 years for which topics $p \in \{k, l\}$ and $q \in \{1,2, \dots ,93\}\backslash\{k, l\}$ are prominent together. Then, we calculate the cosine similarity of the vectors v_k and v_l: larger values indicate that topics k and l were previously combined with similar same topics in the past. The degree of novelty attributable to the topic pair kl is then one minus that cosine similarity. The novelty of the whole case-report is the sum of the novelties due to all new topic pairs. Given the lookback period of 30 years, and in view of the small number of reports in the first part of the 16th century, we report novelty results for the period 1540-1764, for 14,493 case-reports.

At the case-report level, novelty scores range from 0 to 9. There is a positive novelty score in 26.5% of case-reports, but 18.8% are marginally novel, having a score of less than one. Only 1% of novel cases (0.26% of the 14,493 reports) have a novelty score above 5. The cases with the highest novelty scores tend to be in the earliest years of the dataset, but the number of novel cases in the earlier years is much smaller than the number in later decades.[5]

Our case-report-level novelty score consists of two elements: how many novel pairs are present in the case-report and how much the score is discounted because the novel pairs were comparatively easy to create. The discount averages 50% across all cases with novel pairs. Notably, the discount decreases with the number of new topic pairs in a case, suggesting that more novel cases are in fact ones with more novel pairs that are non-obvious. Interestingly, the discount is also strongly negatively correlated with time: earlier cases produce less obvious novel pairs than later ones.

One objection to our implementation of novelty could arise from perusing Figure A.1 in the Appendix, showing the distribution of reports over time. In the data, there is an imbalance in the number of reports over time, with later years having more case-reports. One might therefore conclude that this leads to a bias whereby too much novelty is attributed to the early periods. However, there are two counterbalancing effects that arise due to the temporal imbalance of case reports. On the one hand, the existence of fewer past reports indeed implies a greater chance that an early report in our data has a novel topic pair. But at the same time, the existence of more current reports implies a greater chance that there will be a novel report among them, which given our data (Figure A.1) implies, all else equal, more novel reports in later years than in earlier ones. Thus, from the theoretical standpoint, the net effect of the distribution of cases on estimated novelty is generally ambiguous and no presumption of bias can be maintained.[6]

There are other reasons to doubt that the temporal imbalance in the number of cases inherently biases our novelty calculations. Very importantly, caselaw is cumulative: there was less law and fewer reports in earlier periods than later ones. Therefore, in earlier times, it might have been easier

[5] The mean novelty score for case reports with a positive novelty score is 1.54 before 1600 and 0.59 after 1700. There are 644 case reports with a positive novelty score before 1600 and 1,294 after 1700.

[6] Some pertinent empirical evidence has arisen within debates about whether progress in modern science is slowing. For example, Park et al. (2023: 140) find a constant absolute number of highly disruptive papers over time, despite an exponential increase in the number of papers. This indicates that measured novelty does not necessarily increase as the number of pertinent documents increases.

to find gaps in the law, and thus easier to create novel cases. It is thus not surprising that in our data positive novelty appears in 37% of all cases before 1650, in contrast to only 23% of all cases after 1650. This is not an artifact of data and process, but rather a reflection of cumulative legal development when new caselaw was being created. Moreover, throughout the remainder of the paper, we present empirical evidence that would counter any lingering doubts about whether our estimates of novelty are merely an artifact of the temporal distribution of reported cases. For example, we will show that there are years of high novelty and low novelty in all time periods.

5. Validation

Construction of a new method of measuring legal development requires validation. We do this in five ways. First, we examine summary statistics implied by our novelty measure for broad conformity with general ideas about the process of novelty. Second, we examine a sample of cases showing the congruity in each case between the novelty score, the new topic pairs in the case, and the words of the case themselves. Third, we compare some basic statistics on novelty to those on case citations made in three important works on English legal history: Blackstone's Commentaries (Blackstone 1893 [1765]), written exactly when our data end; Holdsworth (1938 [1903-1938]), produced half-way between the end of our data and today; and the very recent Baker (2019). Fourth, we consider variations in novelty over time, across courts, and between areas of law, examining broad conformity with general knowledge about early-modern English political and legal history. Fifth, we analyze relationships between novelty scores and case citations, examining whether these relationships conform to what might be predicted from the characteristics of novelty and the historical record. The first three exercises appear immediately below. The fourth and fifth appear in the remaining sections of the paper: they are a natural product of the analyses that follow the construction of our measure.

5.1 Broad characterizations of the data

We first compare summary data to that produced in papers focusing on computational measures of novelty in non-legal domains. In our sample, 73.5% of reports have no novel combinations, 16.2% have exactly one novel combination, 5.8% exactly two, and 4.5% have three or more novel combinations. We also find that novelty is highly skewed: 70% of total novelty is in the 10% of cases with the highest novelty scores.[7]

Uzzi et al. (2013) find novel reference combinations in academic articles and define a set of 40% of articles as that set where "novelty may reside". This would correspond roughly to our 26.5% of cases that have at least one novel combination. In Wang et al. (2017: 1419), 89% of all papers make no new combinations of referenced journals. This is more than our percentage of reports without any novel combinations (73.5%). The 11% of papers in Wang et al. (2017) that have one novel combination would correspond roughly to our 10.3% reports with more than one novel combination. Kelly et al. (2021), using a measure of textual similarity, define a breakthrough patent as one that falls in the top 10% of their measure, corresponding again to our 10.3% of case-

[7] Pearson's first coefficient of skewness for the novelty scores is 3.14.

reports that have two or more novel combinations. Kaplan and Vakili (2014) combine expert judgments and topic modeling to characterize 7% of patents as subject-defining, a criterion considerably stronger than novelty. This set would roughly correspond to our 4.4% of case-reports that have novelty scores greater than 2. The distribution of measures of novelty, or their analogs, in all of these papers is highly skewed, much like the distribution of novelty across our case-reports.

Thus, although we generally find more novel cases than studies on patents and academic publications, the general characteristics of our results are in rough conformity with those studies from modern times. These comparisons, on very different datasets, lend credibility to our methodological approach.

5.2 Three cases

We examine three cases spread across the era that we study. The Chamberlain of London's Case (5 Coke 62) from 1590 considered rules on when towns or public corporations were empowered to make ordinances regulating trade. The context was the City of London regulating the sale of cloth. The case prominently featured two topics that had not appeared prominently together in previous decades—Regulating Commerce and Restraints on Trade. The court discussed the role of legislation, custom, the profit motive, and the public welfare in setting limits on regulations on trade. The parallel between our identified new topic pair and the discussion in the case-report is transparent. This case earned a novelty score of 0.82 and was heavily cited in later reports. It has become a classic reference when discussing Coke's fostering of economic liberalism (Smith 2014: 167).

In Radford v. Bludworth from 1661 (1 Lev. 13), a person who became bankrupt on two different occasions sold land during the same time period. Later the commissioners of bankruptcy sold the same land. The question arose as to which sale was valid, with the plaintiff and defendant in the pertinent suit being lessees of the two purchasers of the land. This involved deliberations on timing elements implicitly contained in an earlier statute on bankruptcy. The case resulted in a rule as to how the statute should be interpreted, making explicit reference to how many years must pass between a declaration of bankruptcy and a valid sale. The case prominently featured four topics—Wrongful Possession, Bankruptcy, Rulings on the Calendar, and Clarifying Legislative Acts, making six topic pairs, five of which were new. The parallels between our identified new topic pairs and the case's context are evident. This case earned a 3.17 novelty score.

In Arthur v Bockenham (Fitzgibbon 233) from 1731, a question arose as to the interpretation of the Statute of Wills (1540). At issue was the bequest of a use (a property right), where the will was made before the testator owned the pertinent property right. The case prominently featured two topics that had not been paired prominently in the same case in previous decades—Uses and Clarifying Legislative Acts. Basing their ruling on the construction of the Statute, the judges held that a bequest was not valid if the testator did not own the use at the time when the will was made. The bequest remains invalid even if ownership was procured after the will was made and before

death. The parallel between our identified new topic pairs and the case's substance is clear. This case scored 0.59 on the novelty index and was cited in later reports.

In sum, these three examples show, first, that our novelty measure can identify cases that contribute new elements to caselaw; and, second, that our identification of novel topic pairs coincides with the recombinant context discussed in the case-reports.

5.3 Three legal anthologies

Blackstone's Commentaries (1893 [1765]) are widely regarded as the most complete, and influential, treatise on English law written in the 18th century. Blackstone (1893: 36-37), paraphrasing Fortescue, had the objective of not examining "the critical niceties of the law…[but rather tracing]…the principles and grounds of the law, even to their original elements." Thus, Blackstone provides citations to the core elements that created the law, that is the set of novel sources. We use existing data from Murrell and Schmidt (2011), who extracted and dated all Blackstone's references to the ER from 1559 to 1714.

For the period 1559-1714, our measure of novelty has 2,684 ER cases that have a positive novelty score, while for the same period Blackstone cites a very similar number—2,534 ER cases. Dividing 1559-1714 into chronological halves, at 1636, the earlier period features 38% of our novel cases and 43% of Blackstone's. The correlation of the number of Blackstone cases cited in a year and our number of novel cases in a year is 0.56 ($p<0.001$) for the whole time period, 0.25 ($p=0.025$) for the earlier period, and 0.64 ($p<0.001$) for the later period.

Similarly, Holdsworth's (1938) History of English Law is widely regarded as the most complete, and influential, treatise on English law completed between Blackstone and the mid-20th century. For Holdsworth, reported cases are primary within the "chief authorities for the history of the law at different periods" (Holdsworth (1938 I: xliv; II: vii). For our full sample period, 1540-1765, we find 3,848 ER cases that have a positive novelty score, while Holdsworth cites 1,818 unique cases. Dividing 1540-1765 into chronological halves, at 1652, the earlier period features 33% of our novel cases and 37% of Holdsworth's cited cases. The correlation between the number of Holdsworth cites in a year and our number of novel cases is 0.50 ($p<0.001$) for the whole time period, 0.33 ($p<0.001$) for the earlier period, and 0.42 ($p<0.001$) for the later period.

The same exercise can be carried out for a modern "elementary textbook" (Baker 2019: v), whose goals are naturally more limited than those of Blackstone and Holdsworth: a "concentration…on the main characteristics, institutions, and doctrines of English law". For our full sample period, 1540-1765, Baker cites 731 cases (in contrast to our 3,848 novel cases). This difference is understandable given the contrast between our goals and those of Baker. Dividing 1540-1765 into chronological halves, at 1652, the earlier period features 33% of our novel cases and 70% of Baker's cited cases. The correlation between the number of Baker cites in a year and our number of novel cases is 0.18 ($p=0.006$) for the whole time period, 0.49 ($p<0.001$) for the earlier period and 0.52 ($p<0.001$) for the later period.

Appendix B provides further information. Figure B.1a shows timelines of the yearly percentages of the Blackstone data and the data on the number of our novel cases, while Figure B.1b shows cumulative sums. Figures B.2a, B.2b, B.3a, and B.3b show the analogous information for Holdsworth and Baker. There is a remarkable degree of conformity between all four series, with Blackstone's and Baker's cited cases slightly more concentrated in the earlier period than are our novel cases, while the cumulative timeline for Holdsworth and our data are remarkably close. Thus, comparisons with three standard works in English legal history, produced at very different times, provide strong validation of our methodology. Moreover, this exercise should dispel any lingering doubts about our previous rejection in Section 4 of the presumption that our results have an inherent bias that attributes too much novelty to earlier periods.

Therefore, in comparison with Blackstone's, Holdsworth's, and Baker's citation-based measures, our novelty measure captures distinct, but nevertheless related aspects of legal development. In particular, nothing in the above comparisons casts doubt on the validity of our measure of novelty. At the same time, our measure provides insights into English legal development that extend beyond those captured by Blackstone's, Holdsworth's, and Baker's citations. For example, our method assigns a numerical score indicating the degree of novelty of each case.

6. The temporal landscape of novelty

Providing the most useful information from 14,493 novelty scores requires aggregating over different categorizations of the case-reports. Let η_{it} be the novelty score for the i^{th} case-report produced in year t. Define yearly aggregate novelty as:

$$\theta_t = \sum_i \eta_{it}. \tag{1}$$

Consistent with our current focus on when change happened, expression (1) captures total novelty in a given year, not mean novelty per case in that year.[8]

Figure 1 displays the resultant timeline for 1540-1764 using a three-year centered moving-average. The dashed vertical lines mark the beginnings of new administrations (reigns and the Interregnum). This figure provides an entirely new perspective on the development of English caselaw, a remark that applies equally to all figures that follow.

The ebb and flow of legal development is starkly depicted here. There are rises in novelty soon after 9 of 12 regime-change years, with the exceptions being the beginnings of the reigns of James II (1685), Anne (1702), and George III (1760). Legal changes due to a new administration

[8] In years when the number of cases in the ER is small, perhaps because of plague or war, law would have changed less. Additionally, those cases that future judges would need to be aware of had a higher probability of appearing in the ER. Again, the number of cases would reflect the amount of innovation. Thus in examining the overall development of law one should not use average novelty: more fecund years would inherently have more cases in our dataset.

plausibly emerge because new slates of judges could be chosen on the advent of a new monarch, notwithstanding the tenure clause of the Act of Settlement (Murrell 2021).

The most notable feature of Figure 1 is how much novelty occurred during the administration of Charles II (1660-1685).[9] While that period spans 11% of the years in our data, 27% of all novelty occurred then.[10] In contrast, only 28% of all novelty occurred in the 35% of years between 1685 and the end of the sample period. As we show below, the last half of the 17th century is when Chancery began to contribute notably to novel cases. It is, moreover, a time period of remarkable judges, including Matthew Hale in the common-law courts and a succession of prominent Lord Chancellors in Chancery.[11] Also important is that a change in perspective on the law had occurred during the Civil War and Interregnum. This would have led to novel perspectives once the normal legal routine returned after 1660.[12] Furthermore, many legal issues, including ongoing cases, had their resolutions delayed during the disruptions of the Civil War and Interregnum, concentrating novelty into a smaller time period.[13]

Another fecund period runs from 1583 to 1597: during these 7% of the sampled years, 12% of novelty occurred. It would be tempting to theorize that this rapid development was due to the appearance of the case-reports of Edward Coke at this time. There are certainly many claims that these were hugely influential (Plucknett 1956, Holdsworth 1966). However, results presented in Section 9 provide more nuance to this story: the sheer volume of Coke's reports might have contributed to this uptick in measured novelty, but Coke's reports themselves were not on average novel.

There are also times when legal development is notably slow. The period of political struggle, 1630-1646, was especially barren, with only 2.5% of novelty occurring in the corresponding 8% of the sampled years. This is hardly surprising. More interesting is the fact that legal development in the part of the 18th century covered by our data is especially slow, notwithstanding, again, the tenure clause of the Act of Settlement.[14] Only 21% of novelty is produced in this period covering 28% of the sampled years. In the last 20 years of the sampled period (1744-1764), 9% of all years under consideration, only 6% of novelty is produced, despite the advent of Lord Mansfield on the

[9] Notably, this timing coincides with the analysis of Henriques and Palma (2023), which focuses on higher-level institutions and dates England's divergence as beginning around the time of the Civil War.

[10] The 27% is derived by summing the novelty scores during the years in this particular period and then dividing by the sum of novelty scores across all years.

[11] See the following section for discussion of the important distinction between the common law and equity, the latter applied in Chancery. One example of a case in this time period is Vachel v Vachel and Lemmon (1669: 1 Chancery Cases 129), which has a novelty score of 1.54. This case clarified the role of equity courts in ruling on ambiguous contingencies in wills.

[12] See Holdsworth (1966: 134): "The great rebellion had introduced both in law and in politics the modern atmosphere and the modern outlook, so that…the reported decisions show the emergence of modern rules and doctrines stated in a modern way."

[13] Consider, for example, the previously discussed case of Radford v. Bludworth in 1661. The first event relevant to the facts of this case occurred in 1643.

[14] This accords with Holdsworth (1966: 160) "The development of the common law during the first half of the 18th century was slow. Its procedure was very technical; and its rules of pleading were tending to become more and more subtle and rigid."

King's Bench, a judge who is reputed as having remade commercial law (Lieberman 1989: 88). It is this kind of evidence, challenging long-held conventional views, that provides a powerful argument for constructing innovative measures of legal development. The remaining sections of the paper present many more facts derived from our measures. These facts present a new way of looking at institutional change during the formative years before Britain's Industrial Revolution.

7. The contributions of common law and equity

We now decompose novelty by using data on which court produced which case-reports. Cases were decided in different upper-level courts, with the crucial distinction being between the primary equity court, Chancery, and the two main common-law courts, King's Bench and Common Pleas.[15]

The distinction between equity and common law is fundamental, particularly because it is conventional to refer to the common law as providing a crucial input into comparative development (La Porta et al. 2008, Nunn 2020). Common law used rigid sets of writs, adversarial procedures, and remedies. It relied on juries and its judges were relatively independent from the Crown, which was legislated in 1701 (Murrell 2001). Equity, originating with the King's conscience (as interpreted by the Lord Chancellor, the head of Chancery), provided a more flexible forum, especially in moderating harsh judgments that resulted from the rigors of common law. It did not use juries, its judges investigated facts, which provided a powerful means of discovery, and its chief judge was one of the Crown's most important ministers. Originally, the decisions of equity were not constrained by inherited rules, and thus precedent and case-reporting were relatively less important. The move to a greater use of precedent in Chancery occurred in the first half of the 17th century, with increased reporting soon thereafter (Veeder 1901b: 110). Nevertheless, in the period covered by this paper, the functioning of Chancery was as close to that of civil-law systems as to the English common-law.

Define η_{ijt} as the novelty score of the i^{th} case-report produced by court j at time t. Define κ_{jt} as the aggregate contribution of court j to novelty at time t. Then, noting that the set of case-reports i indexing the summation varies with court j:

$$\kappa_{jt} = \sum_{i} \eta_{ijt}. \tag{2}$$

κ_{jt} can vary over time for two principal reasons: specific courts contribute more reports or judges within specific courts contribute more novelty. We thus decompose (2) in the following way. Let m_{jt} be the number of reports from court j at time t and $\bar{\kappa}_{jt} = \kappa_{jt}/m_{jt}$ be the mean report-level novelty in court j. Then, $\kappa_{jt} = \bar{\kappa}_{jt} \times m_{jt}$, where $\bar{\kappa}_{jt}$ and m_{jt} provide the two components of aggregate novelty contributed by court j at time t. We also examine the share of novelty within any given year produced by each court: κ_{jt}/θ_t, where θ_t is defined by (1). The time series of the resultant four measures are depicted in Figure 2.

[15] Although there are several other courts in our database, these are of lesser importance in contributing cases and we thus classify them as 'other'.

The major story from Figure 2 is the fluctuation of the relative importance of Chancery and the King's Bench. As the 17th century progressed, an emphasis on precedent-based reasoning increased (Grajzl and Murrell 2021b), but as Appendix Figure C.1 shows, this occurred later in Chancery (1675) than in the King's Bench (1640).[16] A quick comparison of Figure 2b and Figure C.1 reveals that there is a strong correlation between the number of reports in the courts and the importance of the idea of precedent. Therefore, the increasing share of Chancery in production of aggregate novelty (Figure 2d) is, at least in part, due to the increased reporting likely arising from Chancery's delayed acceptance of the idea of precedent. Moreover, the locus of litigation in some substantive areas of law moved gradually from the King's Bench to Chancery, notably in finance (Grajzl and Murrell 2022). Figure 2d certainly supports the view that in the era immediately preceding the industrial revolution, "the chief engine of law reform…was…in the court of Chancery" (Jenks 1938: 211).

There are a number of other notable observations gleaned from Figure 2. It is clear that increases in the mean novelty of reports within the two principal courts are important for the large increase in total novelty occurring in the 1660's (Figure 2a): the nature of legal reasoning really did change after the Interregnum. Later in Charles II's reign, however, the number of reports becomes much more important in accounting for aggregate novelty (Figure 2). Interestingly, the mean novelty of Chancery cases is particularly high beginning in the 1620's (Figure 2a), perhaps as a result of the Earl of Oxford's case (1615: 1 Chancery Reports 1), which declared that that equity was superior to the common law. This heightened mean novelty continues to 1680. Last but not least, the fecundity of development from 1583 to 1597 is largely due to the growth of reporting, in which Edward Coke must have been influential. Baker (2019: 51-52) casts the King's Bench as especially innovative at this time, but Figures 2a and 2b suggest that this innovation was primarily a consequence of an increasing number of cases rather than any rise in the mean novelty of decisions.

8. Novelty across areas of law

We now pursue an analogous decomposition for different areas of law. We aggregate the 93 topics into 9 themes. (The allocation of topics to themes is detailed in Appendix Table D.1.) Then one obtains a 14,493-by-9 document-theme matrix, with typical element ρ_{jit}, the proportion of legal theme j within the i^{th} case-report produced in year t. Now define η_{it} as the novelty score for the i^{th} case-report produced in year t. Then, the aggregate amount of novelty attributable to theme j at time t is:[17]

[16] The figure is constructed using the data from Grajzl and Murrell (2021a).

[17] That is, we do not allocate novelty solely to the topics that were in the novel pairings of topics, but rather to all topics in proportion to their presence in the case. Recall that the topics in new pairings are not defined as the most important focus of a case, but rather prominent topics that are used together in a novel fashion. A new pairing of topics could easily provide novelty in a third area of law, as in Arthur v Bockenham discussed above, where an important contribution to caselaw on inheritance was made even though the new topic pairing combined Uses and Clarifying Legislative Acts.

$$\tau_{jt} = \sum_i \eta_{it}\rho_{jit} \tag{3}$$

and the mean novelty attributable to legal theme j at time t is:

$$\bar{\tau}_{jt} = \frac{\sum_i \eta_{it}\rho_{jit}}{\sum_i \rho_{jit}}. \tag{4}$$

In both (3) and (4), the sums are defined over all case-reports i produced in year t. $\sum_i \rho_{jit}$ in (4) is the total contribution of theme j to reports at time t. Then, $\tau_{jt} = \bar{\tau}_{jt} \times \sum_i \rho_{jit}$, emphasizing the relationship between aggregate novelty attributable to theme j and the two components contributing to it: the mean novelty within the area of law and the prevalence of the area of law in the corpus. We also examine the proportion of total novelty within any given year accounted for by each theme: τ_{jt}/θ_t. The time series of the resultant four measures appear in Figure 3.

Four broad characterizations stand out for the majority of themes. An increase in mean novelty within all legal areas begins in 1640 and continues for about thirty years (Figure 3a). Then there is an increase in the aggregate contribution to novelty in all areas of law that begins in 1660 (Figure 3c). This is a product of the confluence of heightened mean novelty and increases in the reporting on all areas of law (Figure 3b). The zeniths of the yearly aggregate contributions of all areas of law are in the third quarter of the 17th century, when the zeniths of each of the two elements of the production of novelty coincide. Lastly, mean novelty within all legal areas declines as one moves into the 18th century (Figure 3a).

However, specific areas of law exhibit differences from these general patterns. Developments in real and personal property are skewed towards the 16th century (Figures 3b,3c). At that time, the courts were gradually moving away from the vestiges of feudalism, a process that was effectively completed by the Tenures Abolition Act of 1660. The opposite skew appears in procedure, jurisdiction, and sources of law, with precedent being an important element of the last (Figures 3b,3c). Therefore, over time, novelty in the legal system is centered more on the system's own functioning and its place in the country's governance, while there is relatively less emphasis on issues more directly relevant to society outside the courts. Notably, for more modern economic themes, contract and debt, families and inheritance, and markets and organization, there is no large difference in the production of novel law between the earlier years and the later years (Figure 3c). Legal development relevant to economic growth was thus occurring relatively continuously over the whole two centuries.

9. The contributions of individual law reporters and judges to novelty

The literature is replete with laudatory and disdainful comments about reporters of cases and judges in this era, from Holdsworth (1966) and Winfield (1925) identifying a few superstars, to Foss's (1857) opinionation on every judge, to discussions of the quality of the reports (Wallace 1855, Veeder 1901a, 1901b; Baker 2019: 186-197). We assess the contributions of reporters and judges to novelty. Using this standard, we identify some unimpressive heroes and some innovative

unknowns, providing many examples of how the construction of our measure allows one to revisit conventional wisdom.

9.1 Reporters

We first apply the analysis to reporters. At this time, when law reporting was unorganized and relied on the initiative of individuals, the quality of different reporters has been a staple discussion within legal history. To investigate quality in terms of novelty, we estimate variants of the following model, one for each reporter r:

$$\eta_{ijt} = \beta_r R_{irt} + \phi_t + \varepsilon_{ijt}, \tag{5}$$

where η_{ijt} is the novelty score in case-report i produced by reporter j in period t. R_{irt} is a dummy variable equal to 1 if $j = r$ and 0 otherwise. The parameter of interest, β_r, is the contribution of reporter r to novelty relative to all other reporters. ϕ_t is a year t fixed-effect. The inclusion of year fixed effects is necessary because sociolegal and economic conditions vary from year to year, affecting the characteristics of the sets of cases that specific reporters report on. ε_{ijt} is a disturbance term.

In estimating β_r, we focus only on those reporters with more than 100 case reports in our data. Figure 4 provides estimates of β_r with 90% confidence intervals based on robust standard errors, showing reporter names, together with the courts and eras on which they reported.[18] Most of the positive and significant estimates appear where there is increased reporting from a specific court, for example, Chancery in the mid 1660's when it was promulgating new rules (Plucknett 1956: 690-693). Nevertheless, this is not artifact of the data since estimates of β_r capture mean novelty per cases reported.

Results on individual reporters challenge established opinions in the literature. Wallace (1855) evaluates the quality of 23 reporters listed in Figure 4 and our results coincide with his in only 12 cases (for example, reversing our ordering of Reports Temp. Finch and Vernon's Chancery Cases). Winfield (1925: 184) refers to Barnardiston as an incompetent and Coke as a great lawyer. We, in contrast, rank them in the reverse order in terms of novelty, while our results are consistent with Winfield's negative views on Keble. Veeden (1901a: 15) comments positively on the Elizabethan reports of Croke, but our results indicate that those reports exhibit below-average novelty. In contrast, Veeden's (1901a) comments on Croke's reports in the reign of James I are more circumspect while our estimates rate these reports as exhibiting above-average novelty. Our results on Ambler, Vesey Senior, and Carthew agree with Veeden's negative assessments, but disagree with the negative assessment of Levinz and the positive one on Ventris (Veeden 1901a: 17, 19; 1901b: 114). Our estimates are very negative on the novelty produced by Keble, who is given a neutral assessment by Veeden (1901a: 15, 19; 1901b: 114, 115). Our estimates are neutral on Dyer,

[18] 9 of 44 are significant, which is 4 times the number expected if the data were random. We do not use corrections for multiple hypothesis testing because we are not conducting multiple tests of the same hypothesis but rather a series of single tests, one per reporter.

whose reputation as reporter was high (Plucknett 1956: 280; Winfield 1925: 187). While Winfield (1925: 190) and Baker (2019: 194) regard Burrow highly, we rate him as average on novelty.

The most notable result is the estimated negative contribution of the reports of Edward Coke, to this day often hailed as the greatest lawyer in Tudor and Stuart times. Holdsworth (1966: 113) states that "amongst all the many Makers of the English Law, Coke is the most important".[19] Winfield (1925: 188) and Plucknett (1956: 280) each make the comment on Coke that he was so highly estimated by the profession that his work was simply referred to as 'The Reports'. If Figure 4 is any guide, then at least in terms of novelty, the development of English law was not advanced by this most celebrated of lawyers. When we examine Coke's role as judge, this assessment is reinforced.

9.2 Judges

The examination of the contribution of judges to novelty is analogous to that of reporters, with small differences. First, reports are attributed to the whole court rather than a particular judge. In any given court, the chief judge is the most influential judge. Therefore, we attribute the judge effect to the chief judge and trace the relationship between chief judges and novelty. Second, chief judges have far fewer cases produced during their period in office than reporters reported. Therefore, we analyze the contribution of chief judges who can be associated with more than 5 reports. For the same reason, we use 20-year period fixed effects instead of yearly fixed effects. We thus estimate the following model for each chief judge c:

$$\eta_{ijt} = \beta_c R_{ict} + \phi_p + \varepsilon_{ijt}, \tag{6}$$

where η_{ijt} is the novelty score in case-report i produced in the court with chief judge j in period t. R_{ict} is a dummy variable equal to 1 if $j = c$ and 0 otherwise. ϕ_p is a 20-year period fixed effect and all other elements in (6) are analogous to those defined in (5).

We have sufficient information to examine the chief judges in the two common law courts, King's Bench and Common Pleas, and in the equity court, Chancery, with the chief judge in the latter referred to as the Lord Chancellor (or Lord Keeper). Figures 5, 6, and 7 provide the results, following the format of Figure 4 and again showing 90% confidence intervals based on robust standard errors.

In Chancery, five judges stand out, all in the middle years of the 17th century.[20] This is not surprising given the information presented in Section 7 (Figure 2b) on the increase in novelty per case in Chancery. Coventry, with the highest novelty per case in Figure 5, received strong praise from Holdsworth (1938 V: 256) and Foss (1857 vi: 284). Hineage Finch (Lord Nottingham), with the lowest novelty score of the five, has been characterized as "the founder of our modern system of equity" (Holdsworth 1938 V: 344) and Foss (1857 vii: 94) was unstinting in praise of Finch.

[19] In a different context, Holdsworth (1938 IV: 205) states that "The manner in which Coke in his usual fashion had generalized from and embellished his authorities made it possible to cite his authority for such an extension."
[20] During the years of the Civil Wars and the Interregnum, the Chancery was run by a committee. Therefore, we do not report results for this period.

Given our results, perhaps Finch is more of a consolidator who receives some of the praise that should be due to prior innovators. Among these, Cooper (Shaftesbury), Bridgeman, and Hyde (Clarendon) merit no mention from Holdsworth when he considers "the fathers of our modern system of equity" (Holdsworth 1938 V: 227). Our estimates indicate that their production of novelty was high. Similarly, Foss (1857 vii: 63, 79, 133) characterizes Hyde as knowing his own deficiencies, quotes contemporaries of Bridgeman who were highly critical of him, and is negative on Cooper, while Wallace (1855) characterizes Bridgeman negatively as Chancellor. Our estimates reveal that all of these judges were important contributors to novelty in Chancery. In contrast, Wallace (1855: 310) classifies both Cowper and Somers as great men, while we assess their contributions to novelty as null or negative relative to other judges.

For judges in the King's Bench, as for reporters and Chancery judges, our estimates often do not confirm the prevailing wisdom. Holdsworth (1938 IV, V) mentions Bromley, Lyster and Wray only in passing, even though our estimates rank them high on novelty. Foss (1857: v) expresses little opinion on the qualities of these three chief judges. Hardwicke, whom our estimates also rank highly, is referred to by Holdsworth (1966: 180) as a judge who won universal approval. Both Holdsworth (1966: 153) and Foss (1857: vii 386) give very high praise to Holt, whom we estimate as contributing comparatively little to novelty. Our evaluations agree with Foss on Saunders, Robert Hyde, and Pemberton, but find the opposite of Foss's admiration of Popham (Foss 1857 vi: 183; vii: 136, 149, 164).

An especially interesting contrast is between our assessment that Hale contributed much novelty while Coke produced less novelty than the average chief judge. Holdsworth (1966: 134) describes Hale as "the greatest common lawyer since Coke", but also notes that "his influence upon the history of English law has not been so great as that of Coke". Our estimates reverse the relative ranking of these two judges as contributors to legal novelty.

Finally, the ER contains relatively few reports from the Court of Common Pleas. Our estimates in Figure 7 thus only capture results for eight judges. Mean novelty in this court is low (Figure 2a) meaning that most judges are cast as producing less-than-average novelty (Figure 7). Our most negative judgment falls on Hobart, who is characterized very favorably by Wallace (1855: 163, 273) as is Anderson, whom we rank positively. Notably, the results on Coke are in line with those appearing above: the assessment of his novelty does not justify his enormous reputation. These results, together with those on reporters and judges in the King's Bench, imply that our data provide a fundamental reassessment of one of the truly great historical figures.

10. Novelty and citations compared

Citations are commonly used as a measure of the quality of court decisions.[21] But citations do not necessarily measure novelty as a core aspect of legal development. Rather, they measure those decisions that eventually come to be used as standard references. Judges cite past cases to note the

[21] For example, summarizing the literature, Cross and Spriggs (2010) conclude that "Citations are the central metric for assessing the significance of opinions, at least from a legal perspective."

authority underlying a decision. Between an initial novel legal development and citation, non-novel summaries of that legal development might be established as authoritative simply because of the status of the summarizer or the quality of the summary. For example, McMahan et al. (2021) show that, in social-science research, the original innovation does not always get the credit, which often passes to the synthetic work that summarizes the innovation more clearly and persuasively. In this section, we examine where the differences between novelty and citations might lie.

We use a dataset of citations in the ER to earlier ER cases (Murrell 2021). The dataset includes citations made in all cases covered by the ER and thus captures the record from the year of a case's decision to 1865.

Unsurprisingly, we find a strong correlation between the number of cites to cases heard in a year and our yearly novelty score (0.430, $p<0.001$). Figures 8a and 8b further probe the relationship between yearly citations and yearly novelty. Specifically, Figure 8a shows each yearly value of novelty (or citations) as a percentage of the sum of all novelty (or citations) over the period 1540-1765. Figure 8b shows the cumulative sums of novelty and citations over time. The figures provide evidence that legal development as measured by the citation record occurs later than legal development indicated by recombinant novelty. In particular, citations attribute more importance to the early 1600's, exactly when Coke's reports were relatively numerous, while a greater percentage of novelty appears in Elizabethan times. These observations are consistent with the comments above on the effects of authoritative summaries: the citation record will indicate that legal development occurs later than suggested by our novelty measure.

Our main analysis of the relationship between novelty and citations is conducted at the case-report level. Our dependent variable is the count of the number of citations to a case-report. The existing literature suggests that the relation between novelty and the probability or number of cites is non-linear (Strumsky 2015: 1454, Wang 2017: 1420). As explanatory variables, we thus use dummy variables reflecting quartiles of positive novelty scores. As citation scores are characterized by overdispersion, we use a negative binomial model, where the conditional mean of the number of citations, y_i, to case i is:

$$E[y_i|\cdot] = exp\{\beta_0 + \beta_1 q_{1i} + \beta_2 q_{2i} + \beta_3 q_{3i} + \beta_4 q_{4i}\}. \quad (6)$$

β_0 is the regression constant. q_{ki} is a dummy variable equal to 1 if the novelty score of case-report i is in the k^{th} quartile of those cases with positive novelty scores.

We do not include court and reporter fixed effects because, in our context, both the court and the reporter are endogenous to the case-report's substantive content. We also do not control for yearly fixed-effects. While novelty of a case-report likely reflects the socioeconomic conditions in the year of a case, it is not likely that, conditional on novelty, socioeconomic conditions affect case

citations.[22] We base statistical inference on robust standard errors clustered at the level of case reporter.

Figure 9 depicts the point estimates in the form of incidence rate ratios (IRR), where IRR_k corresponding to a β_k for $k > 0$ in (6) equals $exp(\beta_k)$. The value of $(IRR_k - 1) \times 100$ is the percentage change in the expected number of citations when the novelty of case-report i increases from zero novelty to the k^{th} quartile of positive novelty scores. We are thus comparing the effect of various levels of novelty with non-novelty.

Figure 9, panel 9a, shows the results for citations at all times (that is, up to 1865), using 90% confidence intervals. The source of citations is divided into three categories, by any court, by the same court, and by other courts, on the premise that citations across different court systems are much harder to garner. Panels 9b and 9c show the results when we divide citations between those that were made within 80 years and those that were made 80 years later or thereafter.[23]

The estimates imply that novel cases attract more citations, but only after many decades. The third quartile of novelty contributes most to citations, primarily via citations from the same court that produced the original decision. These results are broadly consistent with those in the literature on modern scientific articles that also compare novelty and citations. For example, Wang et al. (2017: 1420) find that novel scientific articles are more likely to be cited and that this effect increases with the degree of novelty. They also find that the effect of novelty on citations is delayed, appearing only after some years.

The fact that low levels of novelty are associated with fewer citations than no novelty seems an anomaly. Perhaps this has the same deep cause as the well-known 'uncanny valley' phenomenon first observed in anime and later widely observed in robotics: a slight deviation from normality makes people more uncomfortable than a larger deviation (Mori 2012 [1970]).[24] Moreover, the inverted-U of the effect of novelty on citations has long been a staple of psychology where the intensity of a stimulus has an optimal level for acceptability, with diminishing returns at high levels (Berlyne 1970). Thus, our estimated inverted-U for the effect of novelty on citations is broadly consistent with ideas from experimental psychology.

11. Reflections

[22] Also, it is possible that the novelty produced in a specific year would affect the perceptions of later judges or reporters of the characteristics of the specific year. In that situation, adding year fixed effects would lead to a downward bias in the estimated relationship between novelty and citations. In some sense, the year is rewarded by past contributions of novelty and it is problematic to use this as a control when estimating the effect of novelty on citations (Angrist and Pischke 2015:215).

[23] The cutoff decision for 80 years was a product of experimenting with different cutoffs. The results for different cutoffs are shown in Appendix Figure E.1.

[24] One recent study of this phenomenon has shown that it involves a region of the brain that produces emotional evaluations when examining stimuli in comparison to past experiences (Rosenthal-von der Pütten et al. 2019). This particular region is activated for many types of stimuli, not only visual ones. The fact that a little novelty is counter-productive for the citation of reports might reflect similar phenomena.

We have used a well-accepted concept within the methodology of science, recombinant novelty, to offer a new quantitative perspective on English legal history. We showed that it is possible to construct theory- and data-driven measures of legal development both in the aggregate and within specific areas of law and courts. Our methodology offers the promise of providing a canonical, quantitative record of legal development. This methodology could be easily applied to any large corpus in which documents were produced at different times. Our measures of legal development provide many new facts about English legal history during that country's formative period of modern development. They also provide information pertinent to judging the validity of many elements of conventional wisdom on early-modern English history.

Among the facts, we highlight the following. Within the time period that we study, the years in which Charles II held power were the most fecund in caselaw development as captured by our novelty measure. The production of novelty in caselaw during this Restoration period proceeded at nearly three times the average rate from 1540 to 1764. Thereafter, development slowed: during the 18th-century years covered by our data, novelty in caselaw was produced at only three-quarters of the rate observed for the whole sample period. Indeed, using our metric, the 64 years in the portion of the 18th century covered by our data contribute only as much legal novelty as the 46 years of Elizabeth's reign. And the production of novelty per year in the 18th century is much slower than in the last two decades of Elizabeth's rule, when novelty production was proceeding at a rate per year that was 50% greater than in our sample as a whole. These data hardly endorse the importance for subsequent legal development that is usually attributed to the enactment of judicial independence under the Act of Settlement (1701).[25] Moreover, our estimates place elements of the development of English legal institutions much earlier than is usual in the economics literature, which has especially emphasized developments after the Glorious Revolution of 1688.

Indeed, as we show and as Holdsworth (1966: 160) remarks, "The development of the common law during the first half of the eighteenth century was slow."[26] But Holdsworth (1966: 160) also contends that "[common law] was developing and expanding less rapidly than the parallel system of equity." Our data clearly show the increasing importance of equity, as administered in Chancery, in the development of English caselaw. The common-law courts, particularly the King's Bench, were all important in the production of novelty until 1660. The balance of contributions shifted in the mid-17th century, with Chancery playing a fundamental role thereafter. These data hardly justify the conventional wisdom on the importance of England's common-law that appears in discussions of comparative legal and economic history (e.g., La Porta et al. 2008, Nunn 2020). A system of law with many civil-law elements, equity, was at least as important as the common law.

[25] Murrell (2021) reviews the pertinent literature.

[26] Our data show that, as England approached the Industrial Revolution, novelty in caselaw became relatively more focused on the functioning of the legal system and on the court's rules and philosophy, a finding consistent with Holdsworth's (1966: 160) observation that, for the common law in the first half of the 18th century, "procedure was very technical; and its rules of pleading were tending to become more and more subtle and rigid."

Our results suggest that other elements of conventional wisdom might be challenged by our estimates. In general, our results on judges and reporters do not back up the literature's assessments of particular individuals. Specifically, of the many individual contributors to English legal development, none loomed larger on the legal stage than the world historical figure of Edward Coke. But our evaluation of his record in generating legal novelty casts doubt on his reputation. Or, stated differently, perhaps Coke was the great consolidator rather than the great innovator. This observation is consistent with our results that compare our measure of novelty to the citation record. We show that the citation record places legal development later than our estimates. This is consistent with a finding from the modern academic literature that the works accumulating citations might not be the original innovative contributions, but rather later syntheses (McMahan and McFarland 2021). It is also consistent with Holdsworth's (1938 IV: 205) judgment that "The manner in which Coke in his usual fashion had generalized from and embellished his authorities made it possible to cite his authority for such an extension."

We have offered one model for the derivation of new stylized facts that relies on a systematic quantitative analysis of text data from early modern England. Future work could re-examine the validity of our findings by using alternative approaches to measuring legal novelty. Variants on the method we use could extend to new triples of topics (Strumsky and Lobo 2015). Or one could characterize each case-report by using the whole vector summarizing topic proportions, and measure the novelty of that vector using cosine similarity as in de Vann et al. (2016). Or the method of Uzzi et al (2013) could be adapted, constructing statistical measures by comparing the realized topic pairs to those that might have occurred had topic pairs been randomly generated. Further removed, legal novelty might be judged in terms of dis-similarity to past cases using only the vocabulary characterizing the cases (Kelly et al. 2021). Or novelty could be measured by finding the distance of a vector characterizing a document from the centroid of the vectors characterizing all documents (Jing et al. 2023). Indeed, with the availability of many implementable algorithms to extract information from texts, one interesting future research project would be to examine the relative properties of different methods of measuring novelty.

References

Angrist, J. D., & Pischke, J.-S. (2015). *Mastering 'Metrics: The Path from Cause to Effect.* Princeton, NJ: Princeton University Press.

Baker, J. (2019). *Introduction to English Legal History.* Oxford, UK: Oxford University Press.

Berlyne, D. E. (1970). Novelty, complexity, and hedonic value. *Perception & Psychophysics, 8*(5), 279-286.

Blackstone, W. (1893). *Commentaries on the Laws of England in Four Books. Notes selected from the editions of Archibold, Christian, Coleridge, Chitty, Stewart, Kerr, and others, Barron Field's Analysis, and Additional Notes, and a Life of the Author by George Sharswood.* (Vol. 1). Philadelphia: J.B. Lippincott Co. Available: https://oll.libertyfund.org/title/sharswood-commentaries-on-the-laws-of-england-in-four-books-vol-1

Cross, F. B., & Spriggs II, J. F. (2010). The most important (and best) Supreme Court opinions and justices. *Emory Law Journal, 60*, 407-502.

De Vaan, M., Stark, D., & Vedres, B. (2015). Game changer: The topology of creativity. *American Journal of Sociology, 120*(4), 1144-1194.

DiMaggio, P., Nag, M., & Blei, D. (2013). Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of U.S. government arts funding. *Poetics, 41*, 570-606.

Foss, E. (1857). *The Judges of England with Sketches of their Lives.* London, UK: Longman.

Fox, D. (1996). Bona fide purchase and the currency of money. *The Cambridge Law Journal, 55*(3), 547-565.

Grajzl, P., & Murrell, P. (2021a). A machine-learning history of English caselaw and legal ideas prior to the Industrial Revolution I: Generating and interpreting the estimates. *Journal of Institutional Economics, 17*(1), 1-19.

Grajzl, P., & Murrell, P. (2021b). A machine-learning history of English caselaw and legal ideas prior to the Industrial Revolution I: Applications. *Journal of Institutional Economics, 17*(2), 201-216.

Grajzl, P., & Murrell, P. (2022). Using topic-modeling in legal history, with an application to pre-industrial English caselaw on finance. *Law and History Review, 40*(2), 189-228.

Grajzl, P., & Murrell, P. (2023). Of families and inheritance: Law and development in England before the Industrial Revolution. *Cliometrica*, *17*(3): 387-432.

Hadamard, J. (1945). *An Essay on the Psychology of Invention in the Mathematical Field.* Newy York, NY: Dover Publications.

Harris, R. (2000). *Industrializing English Law: Entrepreneurship and Business Organization, 1720-1844.* Cambridge, UK: Cambridge University Press.

Henriques, A., & Palma, N. (2023). Comparative European institutions and the little divergence, 1385-1800. *Journal of Economic Growth, 28*, 259-294.

Holdsworth, W. (1938). *A History of English Law.* London, UK: Metheun.

Holdsworth, W. (1966). *Some Makers of English Law.* Cambridge, UK: Cambridge University Press.

Jenks, E. (1938). *A Short History of English Law.* Fifth Edition. London, UK: Methuen & Co. Ltd.

Jing, E., DeDeo, S., & Wright, D. R. (2023). Sameness entices, but novelty enchants in fanfiction online. Retrieved from arXiv:1904.07741v2 [cs.CL] or https://doi.org/10.48550/arXiv.1904.07741

Kaplan, S., & Vakili, K. (2015). The double-edged sword of recombination in breakthrough innovation. *Strategic Management Journal, 36*(10), 1435-1457.

Kelly, B., Papanikolaou, D., Seru, A., & Taddy, M. (2021). Measuring technological innovation over the long run. *American Economic Review: Insights, 3*(3), 303-320.

La Porta, Rafael, Florencio Lopez-de-Silanes, and Andrei Shleifer. (2008). The economic consequences of legal origins. *Journal of Economic Literature*, *46*(2): 285-332.

Lieberman, D. (1989). *The Province of Legislation Determined: Legal Theory in Eighteenth-Century Britain.* Vol. 14. Cambridge, UK: Cambridge University Press.

McMahan, P., & McFarland, D. A. (2021). Creative destruction: the structural consequences of scientific curation. *American Sociological Review, 86*(2), 341-376.

Mori, M. (2012 [1970]). The uncanny valley. Translated by K. F. MacDorman and N. Kageki. *IEEE Robotics & Automation Magazine, June*, 98-100.

Mueller, B. (2022). Recombination for innovation and market impact: Samples and features in hip hop music. *Journal of Evolutionary Economics, 32*, 929–953.

Murrell, P. (2017). Design and evolution in institutional development: The insignificance of the English Bill of Rights. *Journal of Comparative Economics, 45*(1), 36-55.

Murrell, P. (2021). Did the independence of judges reduce legal development in England, 1600-1800? *Journal of Law and Economics, 64*(3), 539-565.

Murrell, P., & Schmidt, M. (2011). The coevolution of culture and institutions in seventeenth-century England. SSRN working paper.

Nelson, R. R., & Winter, S. G. (1985). *An Evolutionary Theory of Economic Change.* Harvard university press.

North, D. C., Wallis, J. J., & Weingast, B. R. (2009). *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History.* New York, NY: Cambridge University Press.

Nunn, N. (2020). The historical roots of economic development. *Science*, *367*(6485), eaaz9986.

Ogilvie, S., & Carus, A. W. (2014). Institutions and economic growth in historical perspective. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of Economic Growth, Vol. 2.* Amsterdam: Elsevier.

Park, M., Leahey, E., & Funk, R. J. (2023). Papers and patents are becoming less disruptive over time. *Nature, 613*, 138–144.

Plucknett, T. F. T. (1956). *Concise History of the Common Law.* Boston: Little, Brown and Co.

Renton, A. W. (1900-1932). *The English reports.* Great Britain. Parliament. House of Lords. Edinburgh, UK: W. Green & Sons.

Roberts, M. E., Stewart, B. M., & Airoldi, E. M. (2016). A model of text for experimentation in the social sciences. *Journal of the American Statistical Association, 111*, 988-1003.

Rosenthal-von der Pütten, A. M., Krämer, N. C., Maderwald, S., Brand, M., & Grabenhorst, F. (2019). Neural mechanisms for accepting and rejecting artificial social partners in the uncanny valley. *Journal of Neuroscience, 39*(33), 6555-6570.

Schumpeter, J. A. (1934). *The Theory of Economic Development.* Cambridge: Harvard University Press.

Smith, D. C. (2014). *Sir Edward Coke and the Reformation of the Laws : Religion, Politics and Jurisprudence, 1578–1616.* Cambridge, UK: Cambridge University Press.

Strumsky, D., & Lobo, J. (2015). Identifying the sources of technological novelty in the process of invention. *Research Policy, 44*(8), 1445-1461.

Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical combinations and scientific impact. *Science, 342*(6157), 468-472.

Veeder, V. V. (1901a). The English Reports, 1292-1865. *Harvard Law Review, 15*(1), 1-25.

Veeder, V. V. (1901b). The English Reports, 1292-1865. II. *Harvard Law Review, 15*(2), 109-117.

Verhoeven, D., Bakker, J., & Veugelers, R. (2016). Measuring technological novelty with patent-based indicators. *Research Policy, 45*(3), 707-723.

Veugelers, R., Wang, J., & Stephan, P. (2022). Do funding agencies select and enable risky research: Evidence from ERC using novelty as a proxy of risk taking. No. w30320. National Bureau of Economic Research.

Wallace, J.W., (1855). *The Reporters: Chronologically Arranged: with Occasional Remarks Upon Their Respective Merits*. Third Edition. Philadelphia: T. & JW Johnson.

Wang, J., Veugelers, R., & Stephan, P. (2017). Bias against novelty in science: A cautionary tale for users of bibliometric indicators. *Research Policy, Volume 46, Issue 8*, 1416-1436.

Whalen, R., Uzzi, B., & Mukherjee, S. (2017). Common law evolution and judicial impact in the Age of Information. *Elon L. Rev., 9*, 115.

Winfield, P. H. (1925). *The Chief Sources of Legal History*. New York: Burt Franklin.

Youn, H., Strumsky, D., Bettencourt, L., & Lobo, J. (2015). Invention as a combinatorial process: evidence from US patents. *Journal of the Royal Society Interface, 12*(106), 20150272.

Figure 1: Aggregate recombinant novelty in English caselaw, 1540-1764

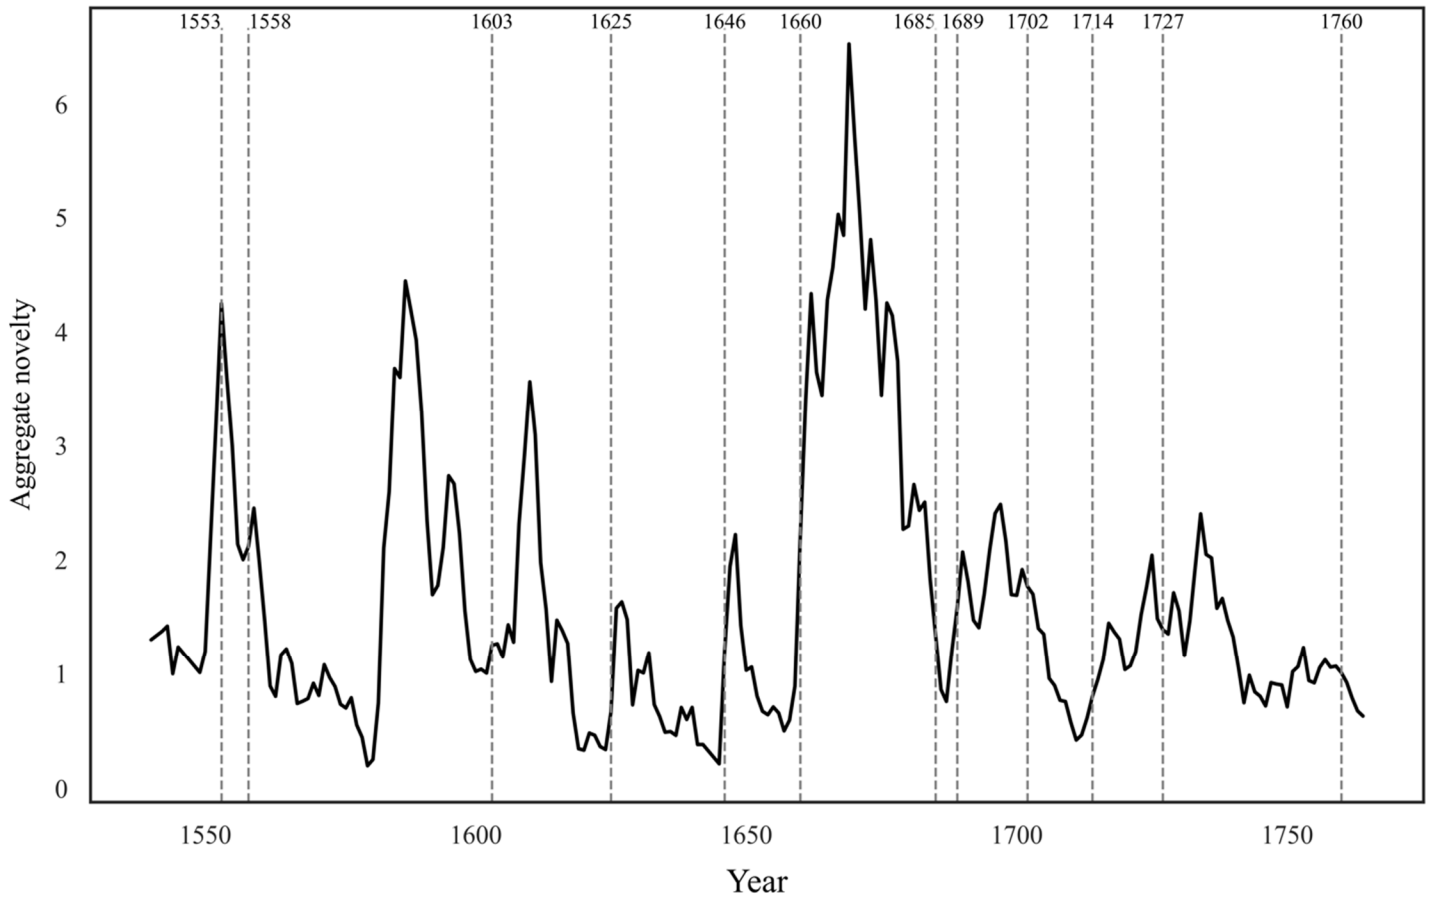

Figure 2: The contributions of courts to novelty, 1540-1764

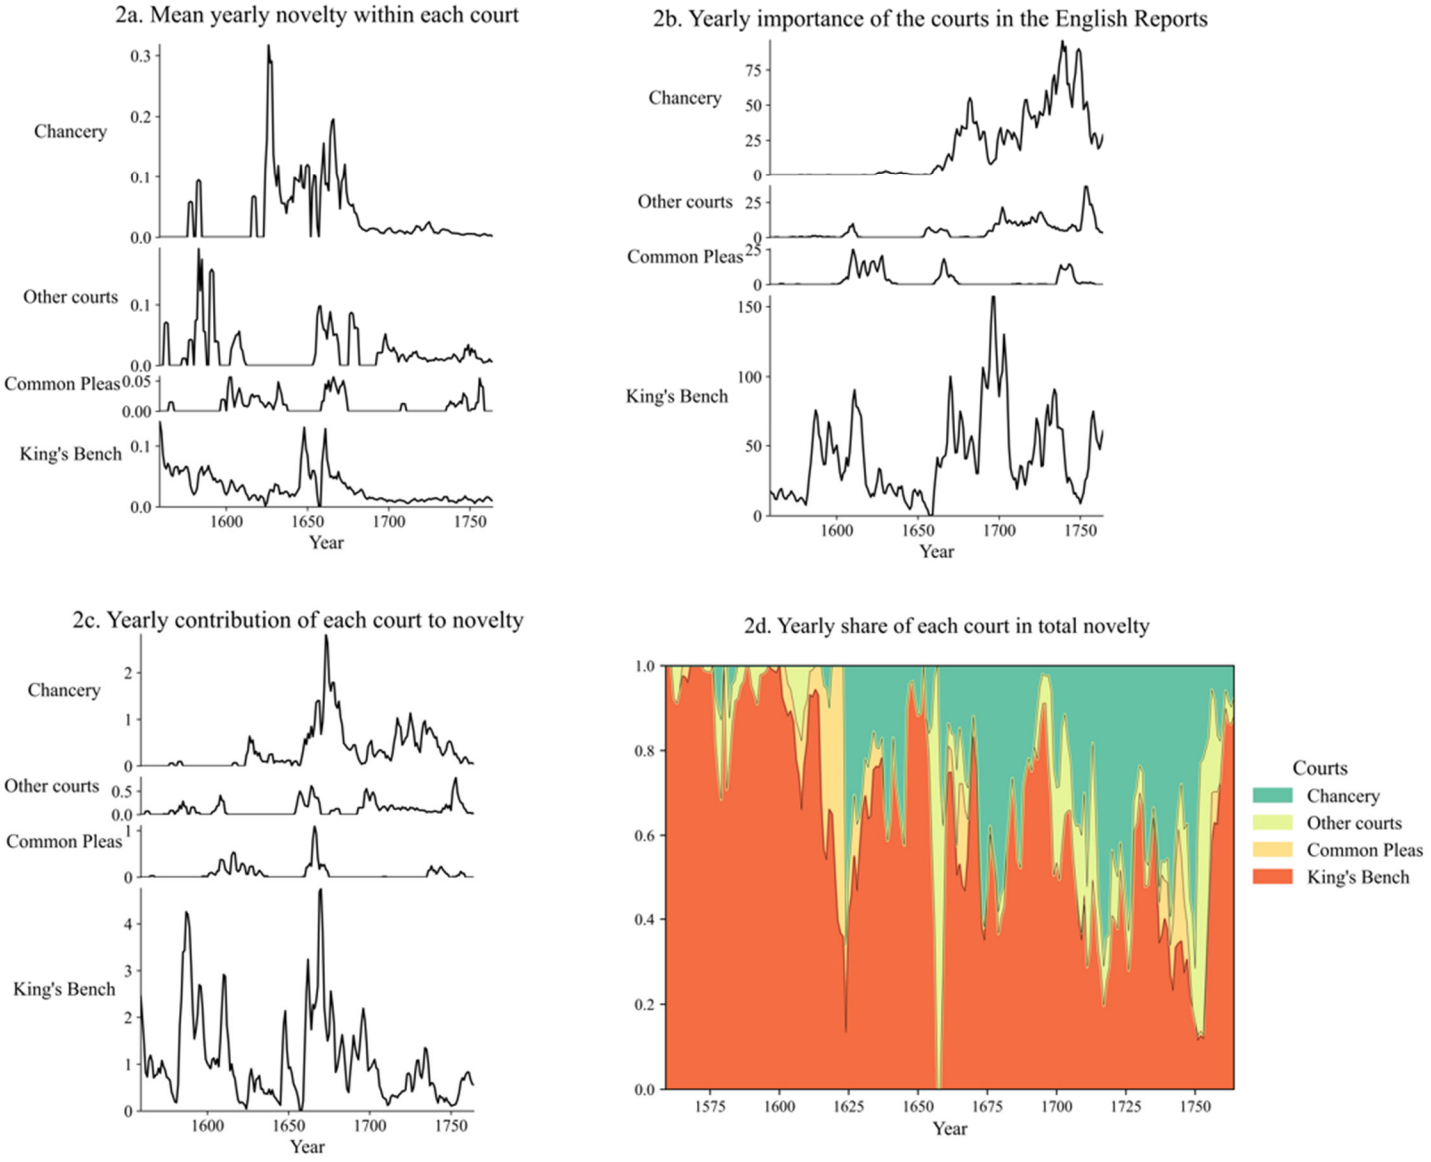

Figure 3: The contributions of areas of caselaw to novelty, 1540-1764

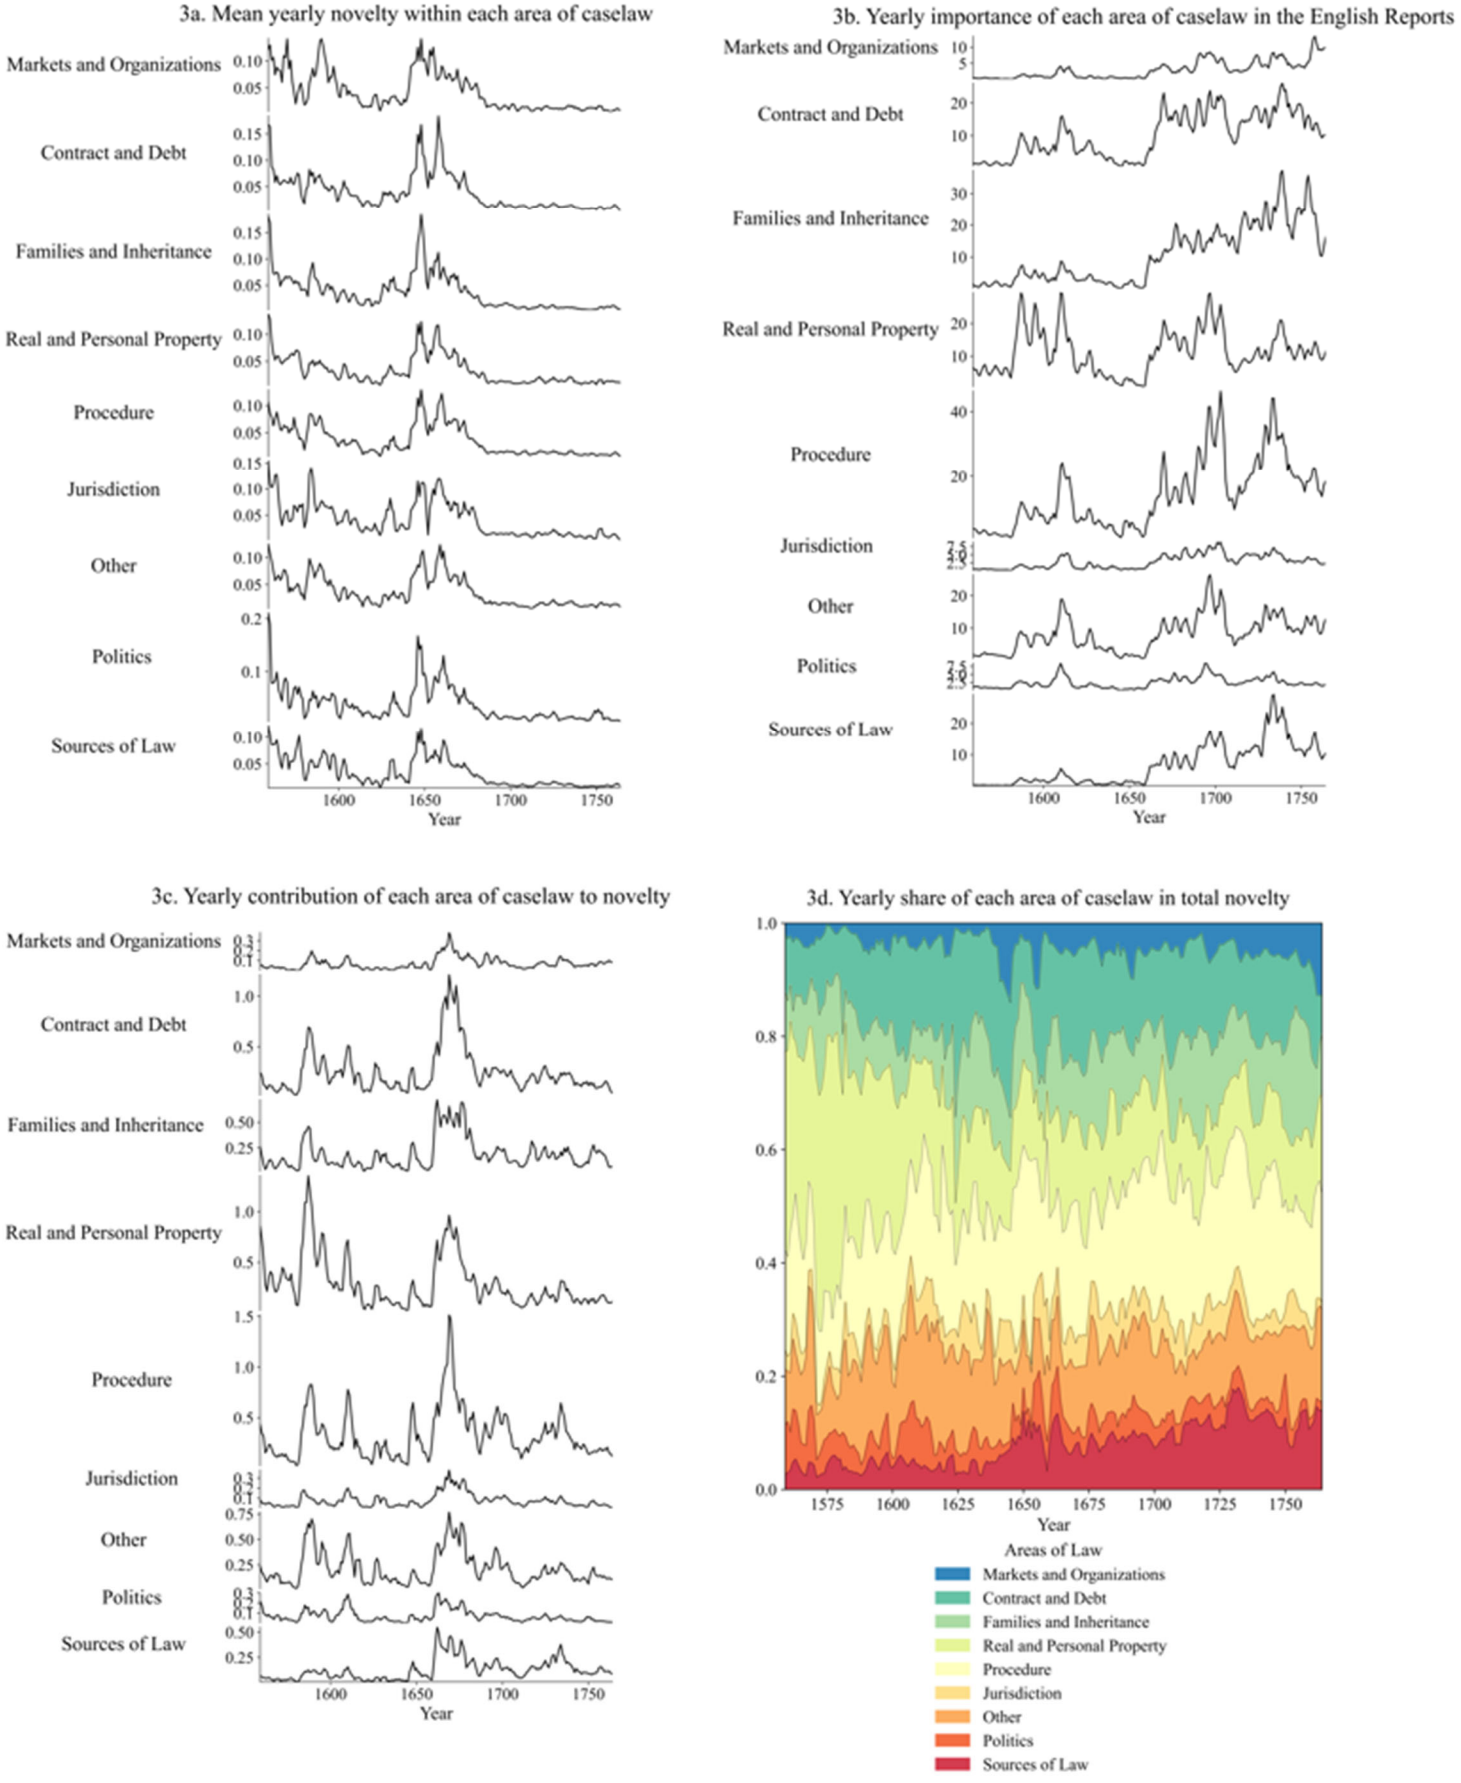

Figure 4: Relationship between reporter and novelty, controlling for year fixed effects

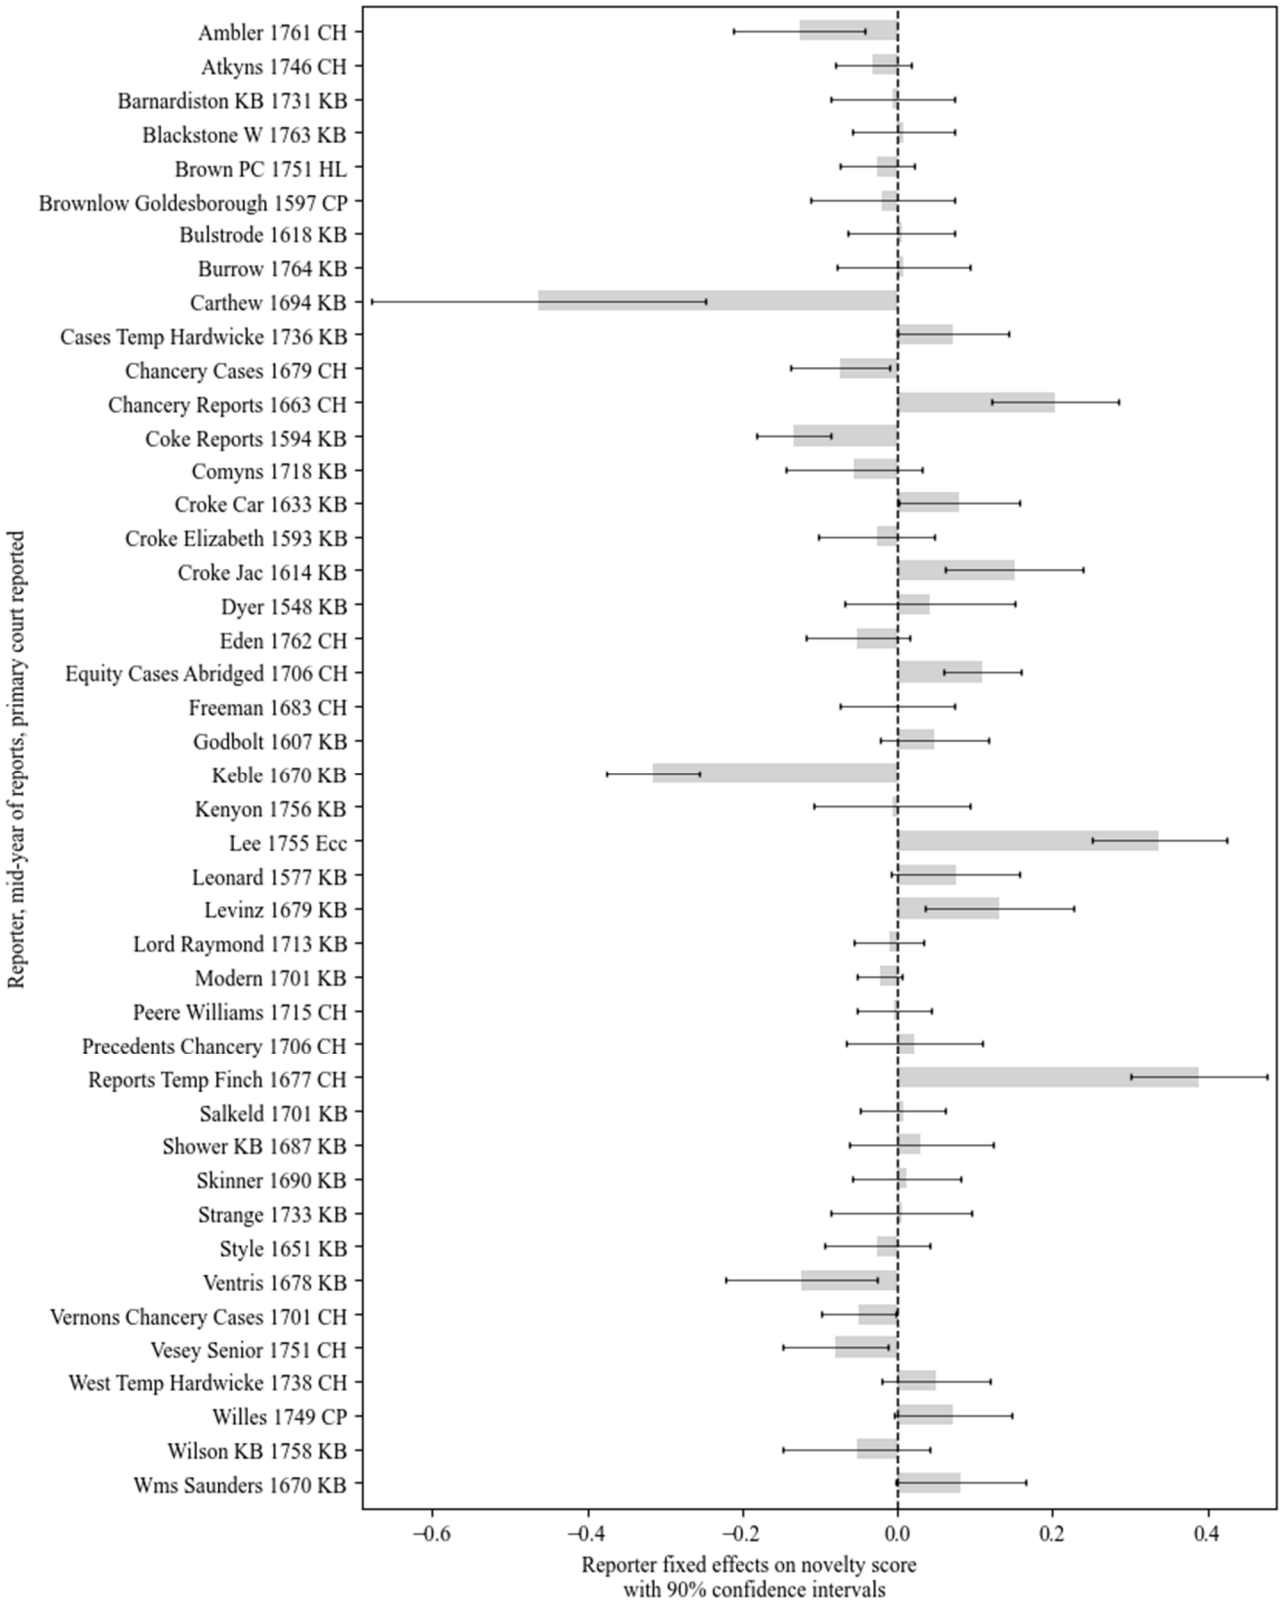

Figure 5: Effect of Lord Chancellor on novelty in Chancery, controlling for period fixed effects

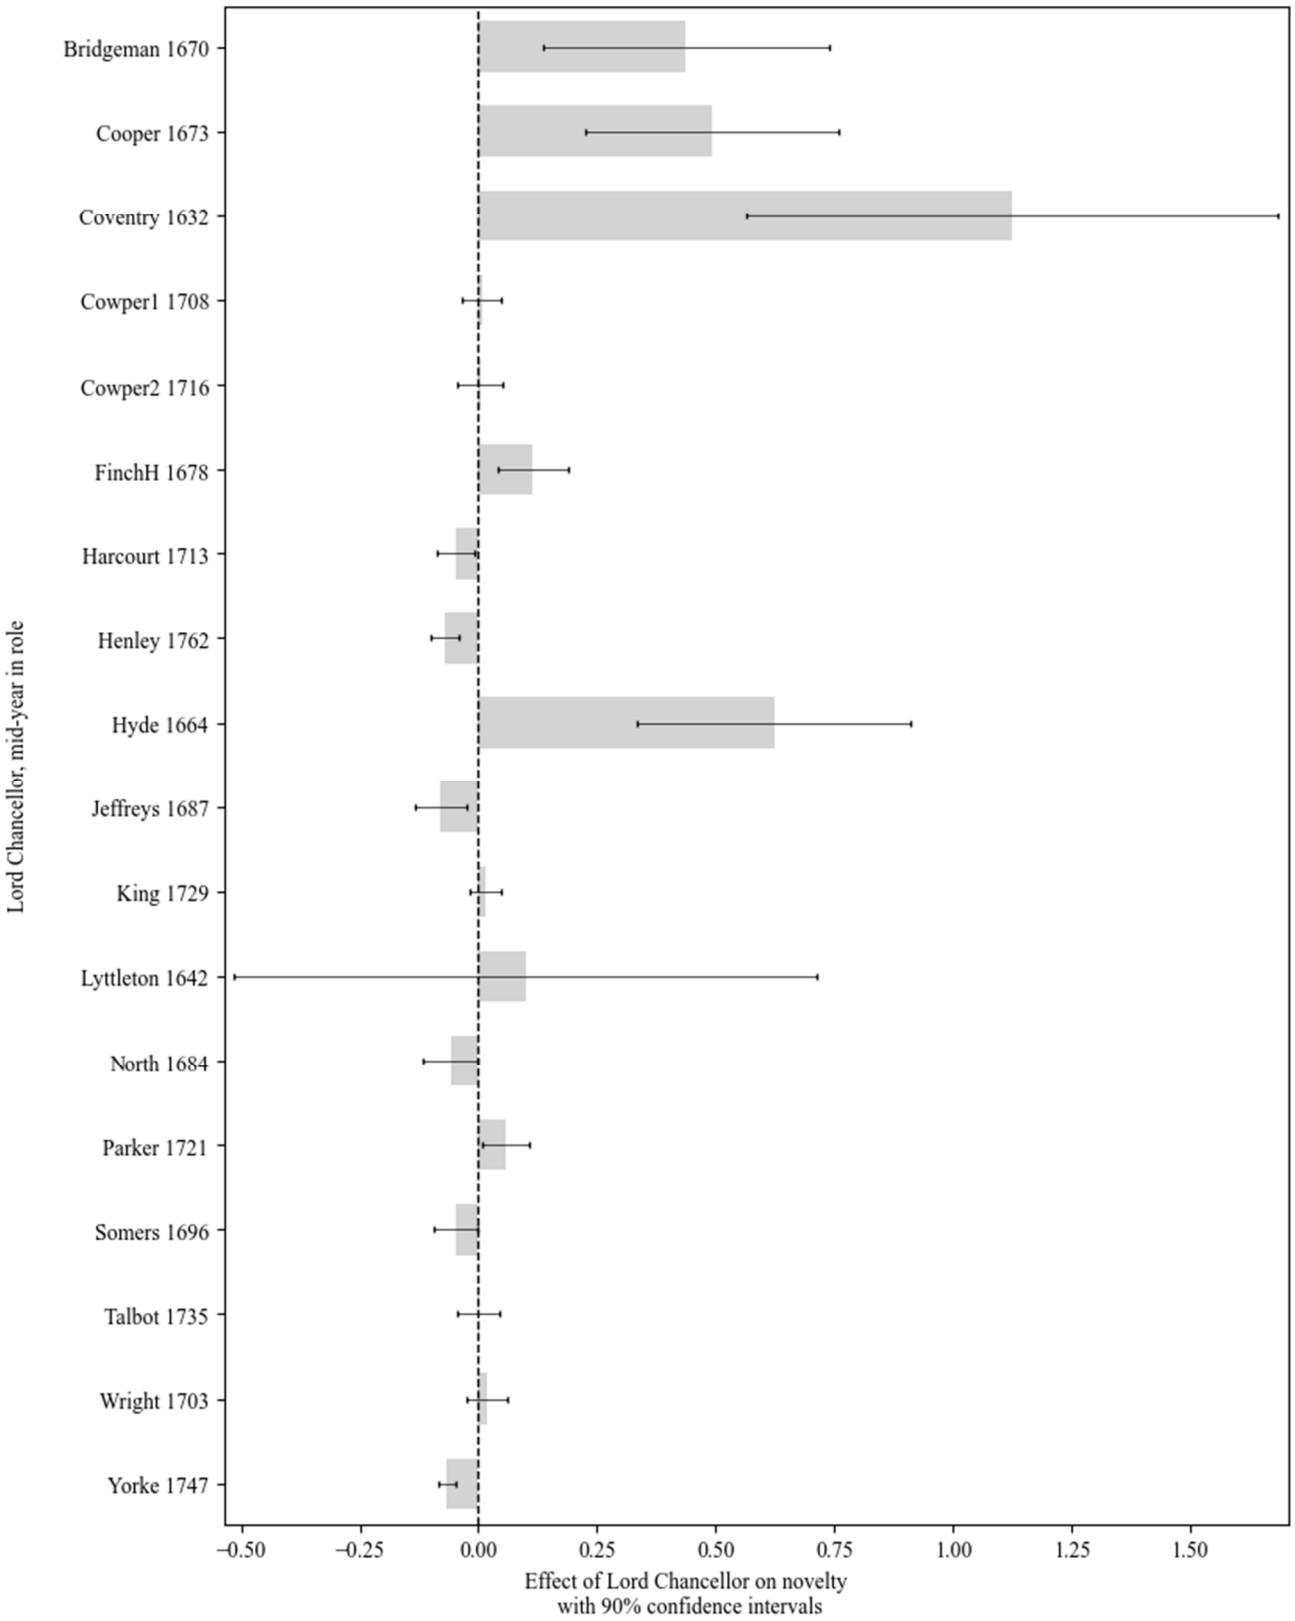

Figure 6: Effect of chief justice of King's Bench on novelty, controlling for period fixed effects

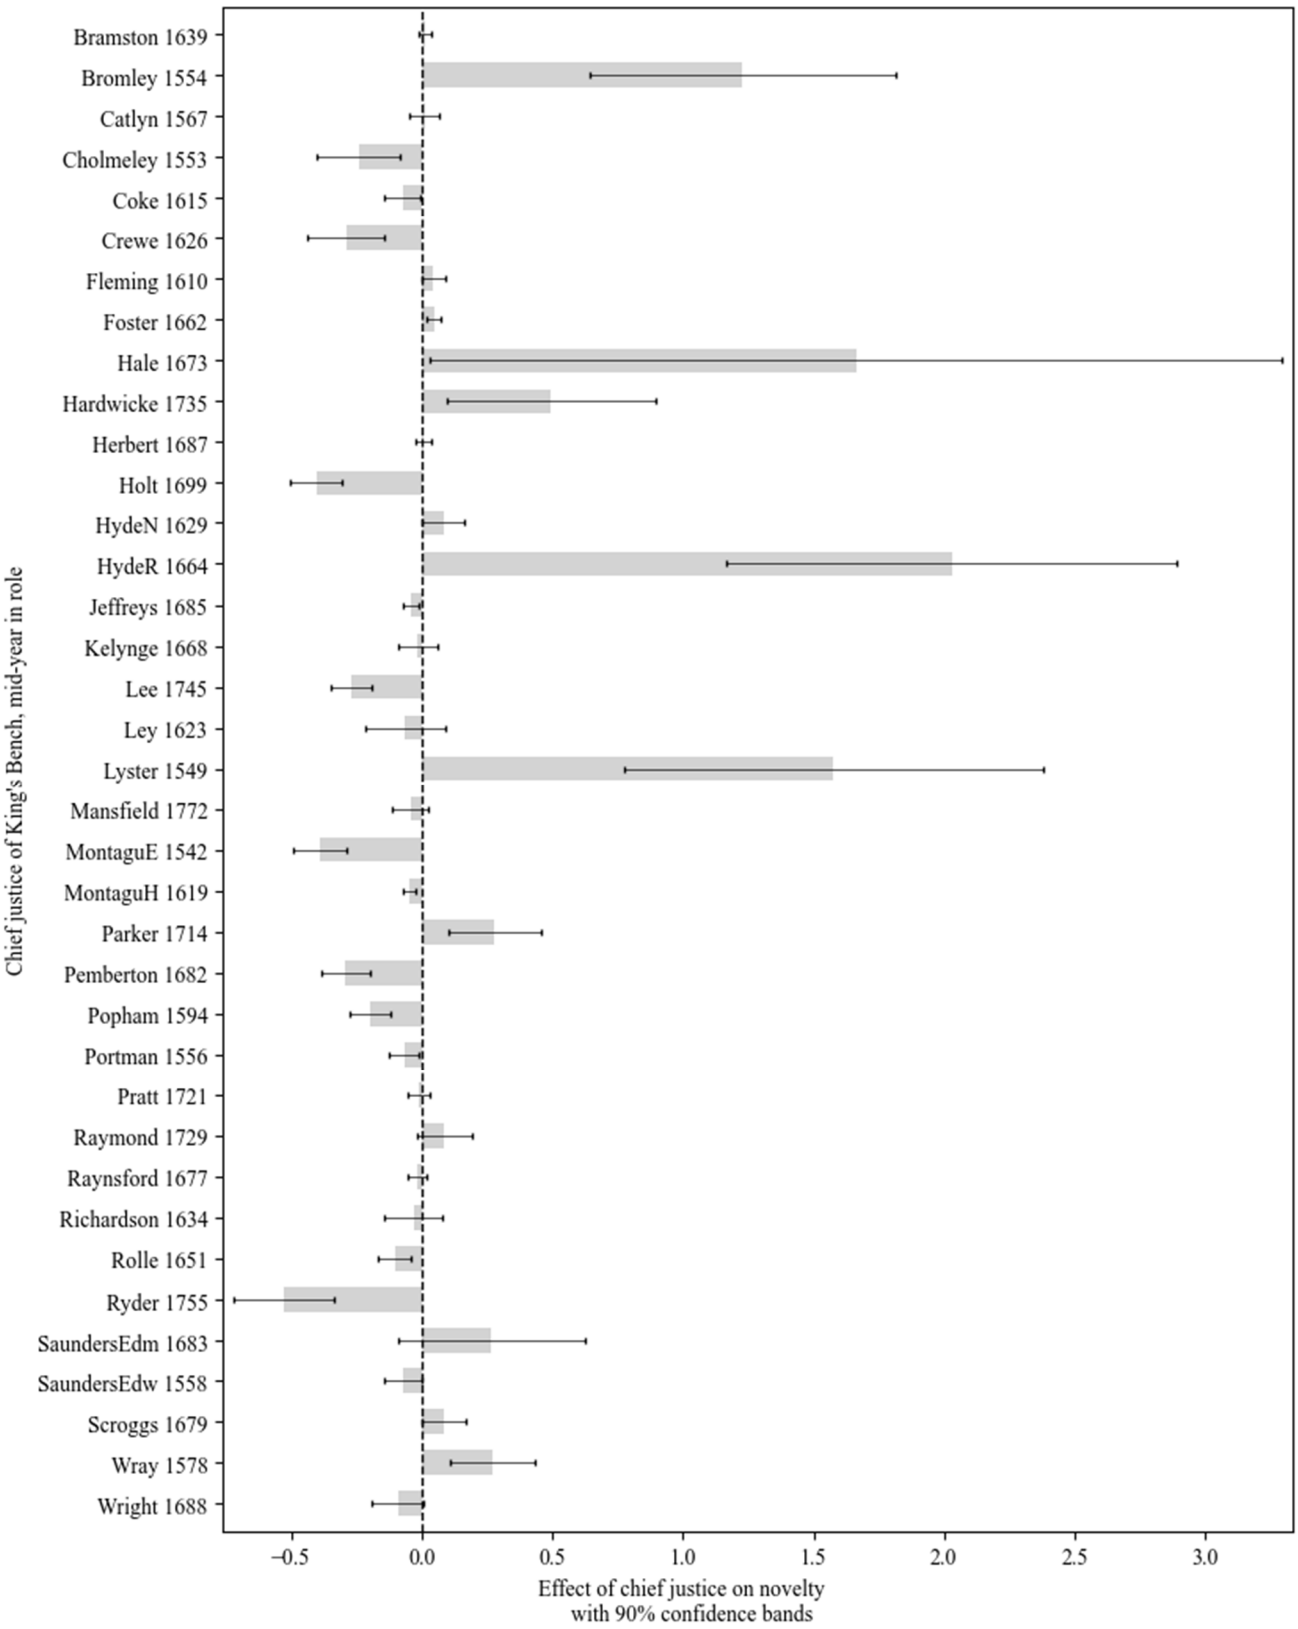

Figure 7: Effect of chief justice of Common Pleas on novelty, controlling for period fixed effects

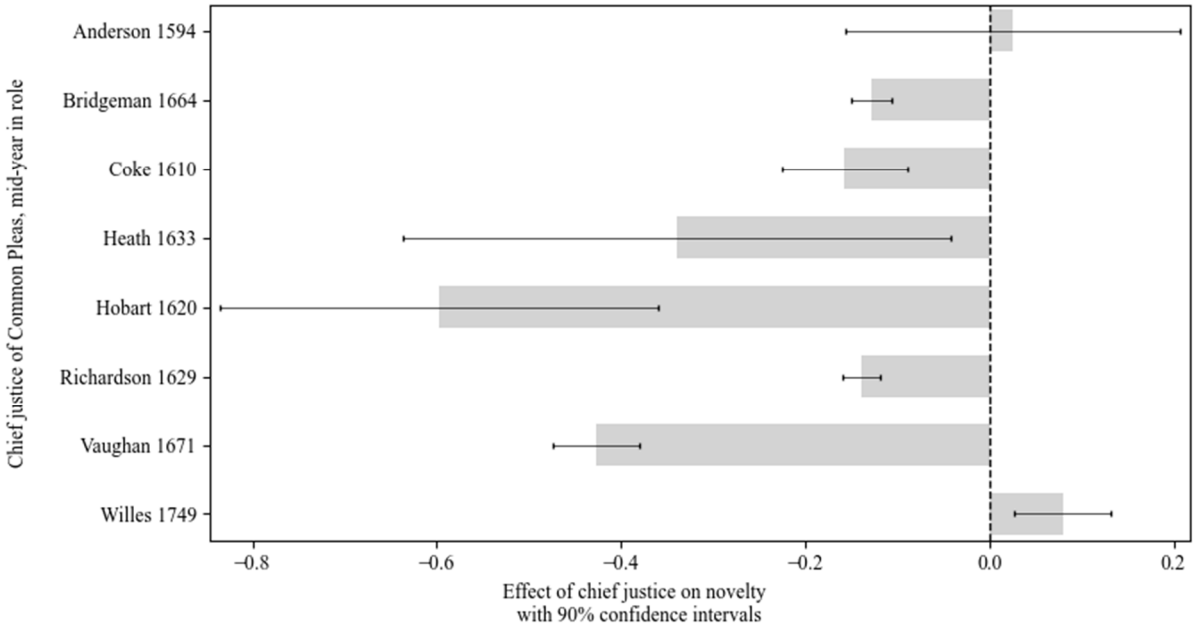

Figure 8a: Comparing yearly novelty and cites

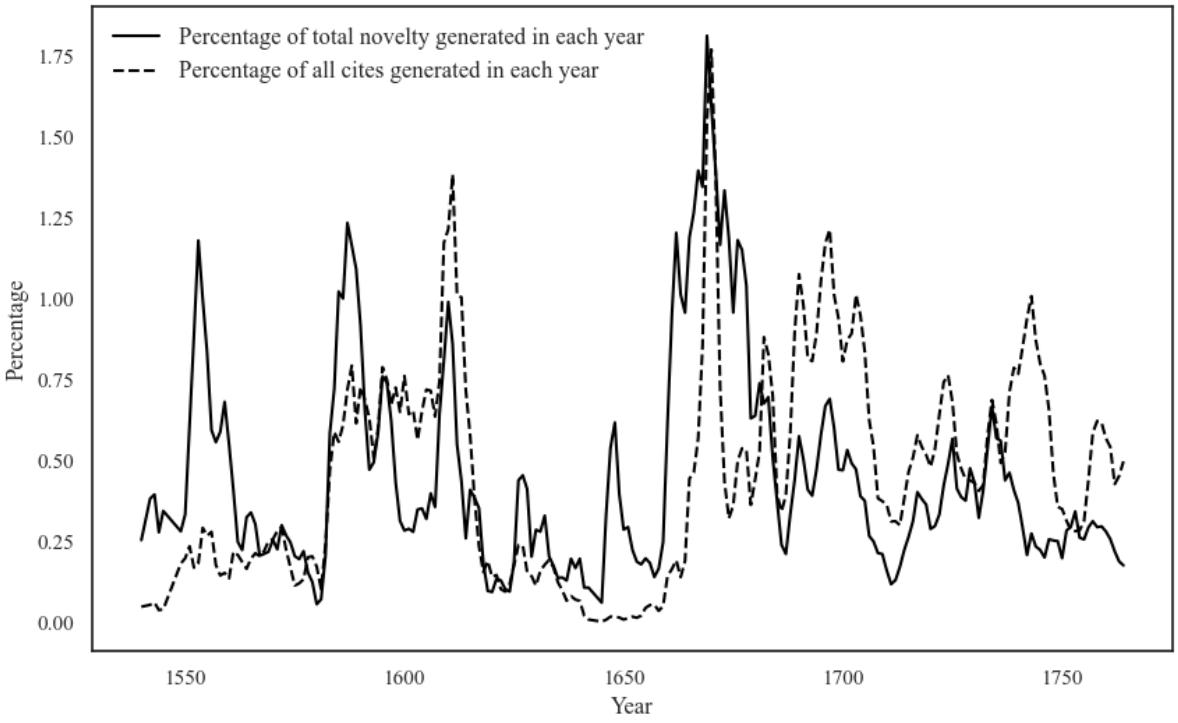

Figure 8b: Cumulative sum of percentage novelty and percentage citations, yearly

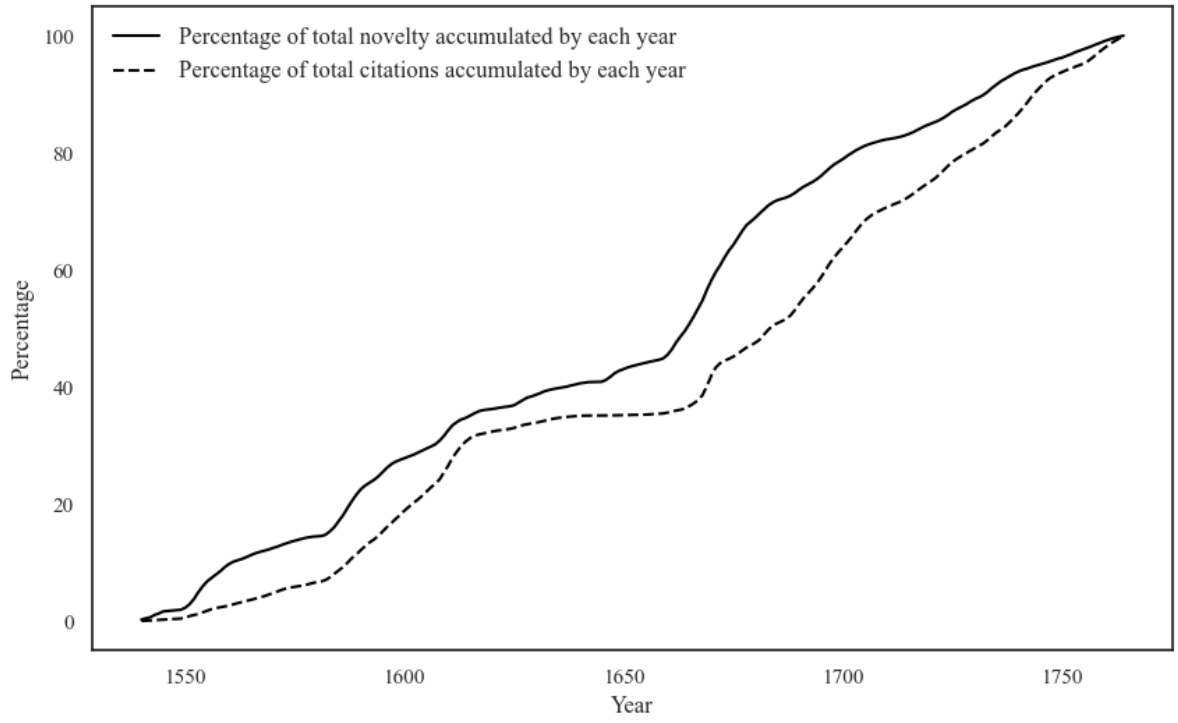

Figure 9: Relationship between novelty and citations,
by novelty quartile and time between cited and citing case

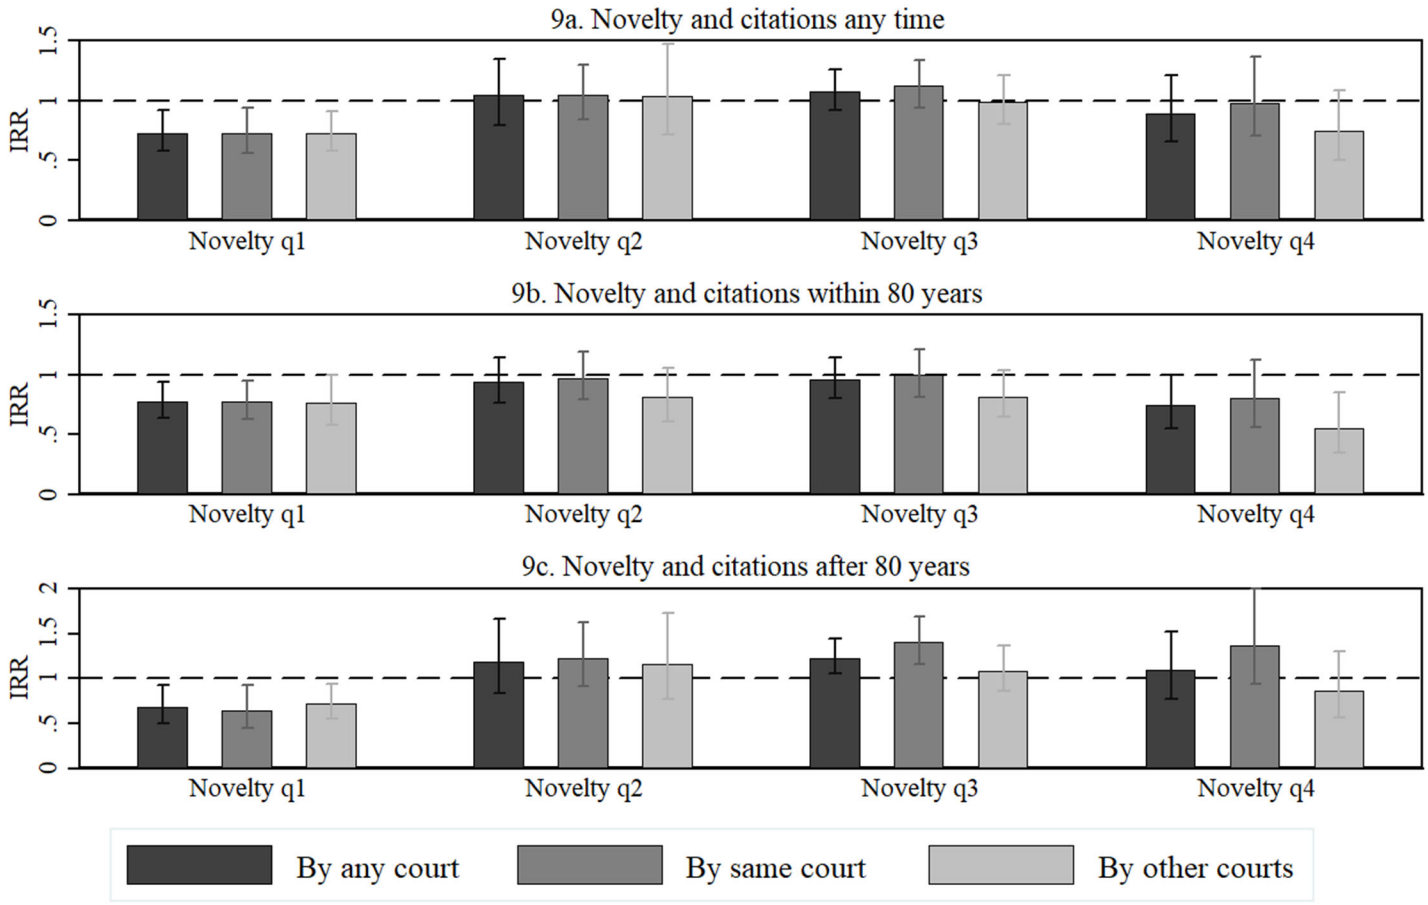

Note: 90% confidence intervals using robust standard errors clustered on reporters

Supplementary Appendixes.

Measuring Institutional Innovation: Recombinant Novelty in Early-Modern English Caselaw

Peter Grajzl Peter Murrell

August 20, 2024

Appendix A: Construction of the dataset

The data from Grajzl and Murrell (2021a) can be found here. There, the R file ERSProcessedData.rda contains an object 'meta', which lists for every document the year, the number of words, the English Reports citation, the court, and the number of citations. These are the items of metadata that are used to produce the results that appear in the paper. They are used in conjunction with the data in the file mtheta.csv, which is the original 52,949×100 document-topic matrix produced by Grajzl and Murrell (2021a). The order of topics in mtheta.csv is the same as the order in Appendix D of the supplementary materials to Grajzl and Murrell (2021a), which means that the topic names can be appropriately assigned to the data on topics. The supplementary materials can be found at here or here.

The original data set contained 52,949 documents. Removing those with less than 500 words leaves 14,549. Removing those from before the year 1501 leaves 14,523. The earliest document is then from the year 1509. We begin our data analysis with that data set. The program sets the lookback period for similar topic pairs at 30 years. That means that we can calculate values for novelty beginning in 1539 and ending in 1764, leaving 14,494 documents. Since the one document in 1539 has a zero novelty score, we discard that document from the presented results, leaving 14,493 observations .

As DiMaggio et al. (2013) argue, one of the valuable features of topic modeling is that non-substantive words are consigned to residual topics, thus improving the quality of the substantive topics. In the Grajzl and Murrell (2021a) dataset there is a topic named Non-Translated Latin. This was excluded from the analysis because its presence or absence does not indicate novelty. Also, there are six topics reflecting various reporting styles, such as Keble Reporting. These were excluded from the topics because changes in reporting style, while indicating novelty in presentation, would not necessarily indicate novelty in substance. These exclusions resulted in the use of a dataset with 93 topics. After the reduction to 93 topics, the document-topic matrix was renormalized so that the sum of topic proportions for each case was 1: all documents were completely apportioned between the 93 topics. A full list of the dropped topics is: Topic 8: Vesey Footnotes; Topic 11: Keble Reporting; Topic 16: Modern Reporting; Topic 26: Non-Translated Latin; Topic 72: Coke's Procedural Rulings; Topic 79: Coke Reporting; and Topic 83: Vesey Reporting.

Figure A.1: The distribution of case reports for which novelty was assessed, 1540-1764

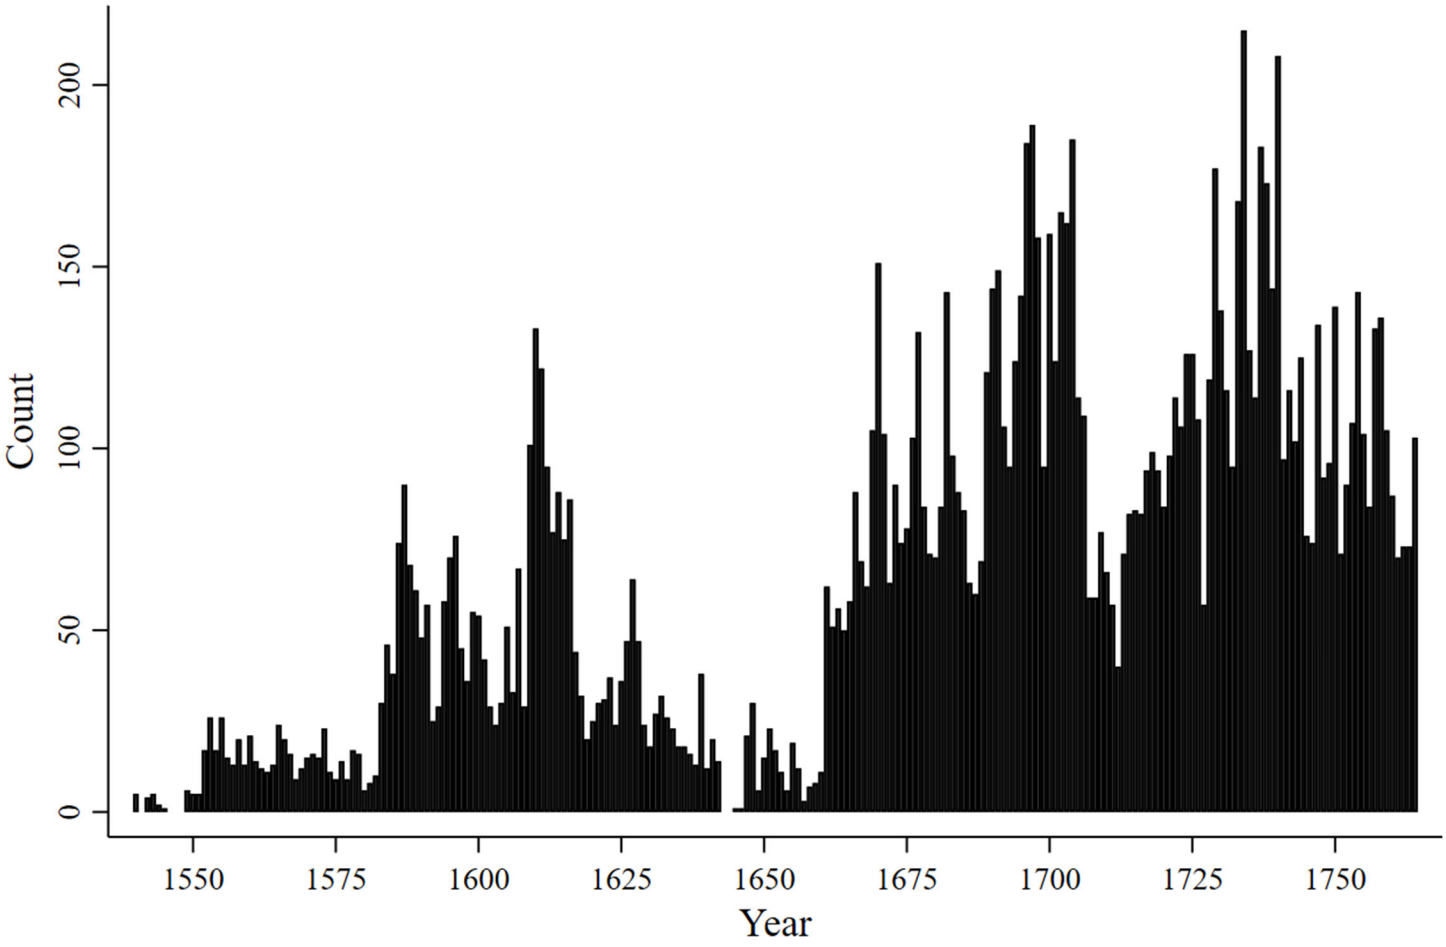

Appendix B: Comparisons between the yearly timelines of cases referenced in Baker (2019) and Blackstone (1893) and the number of novel cases estimated in this paper

The Blackstone data were collected by Murrell and Schmidt (2011) using Python programs applied to the Online Library of Liberty version of Blackstone (1893; https://oll.libertyfund.org/titles/sharswood-commentaries-on-the-laws-of-england-in-four-books-2-vols. All variants of Blackstone's idiosyncratic methods of citation to the English Reports were systematically collected and assigned to the appropriate volume of the English Reports. Years were assigned assuming a linear relationship between the page of the report and the year of the report. The relevant linear function was constructed using the starting year of the first report, the date of the first case, the ending year of the report, and the date of the last case.

The Holdsworth data on citations to the English Reports were scraped and processed using Python programs applied to the 12 OCR'ed Holdsworth volumes downloaded from the Internet Archive (https://archive.org/). To construct a complete set of 12 OCR'ed volumes, volumes from different editions were combined to provide the text input into the computation. The dates of cases were provided by Holdsworth alongside his nominative citations.

Baker (2019) gives the years for his cited cases. Extracting that information from the Kindle version and processing it with a Python program was therefore straightforward.

Figure B.1a: Yearly percentage of novel and Blackstone cases

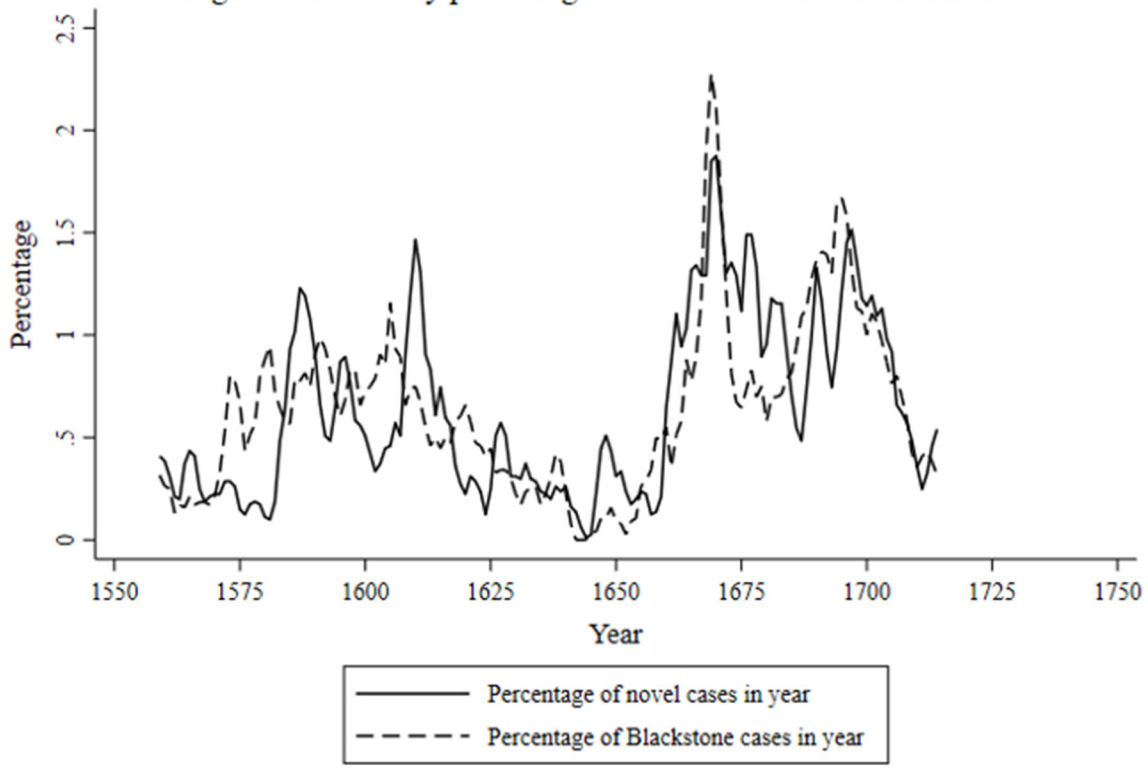

Figure B.1b: Cumulative sum of percentage of novel and Blackstone cases, yearly

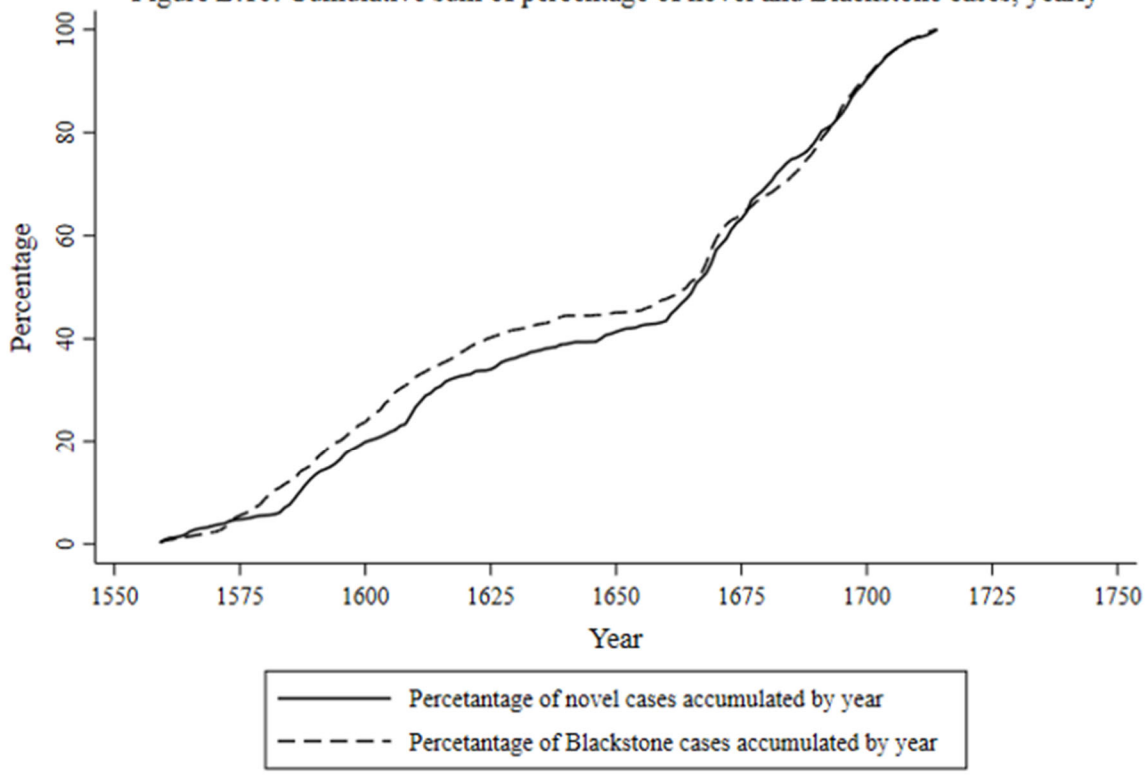

Figure B.2a: Yearly percentage of novel and Holdsworth cases

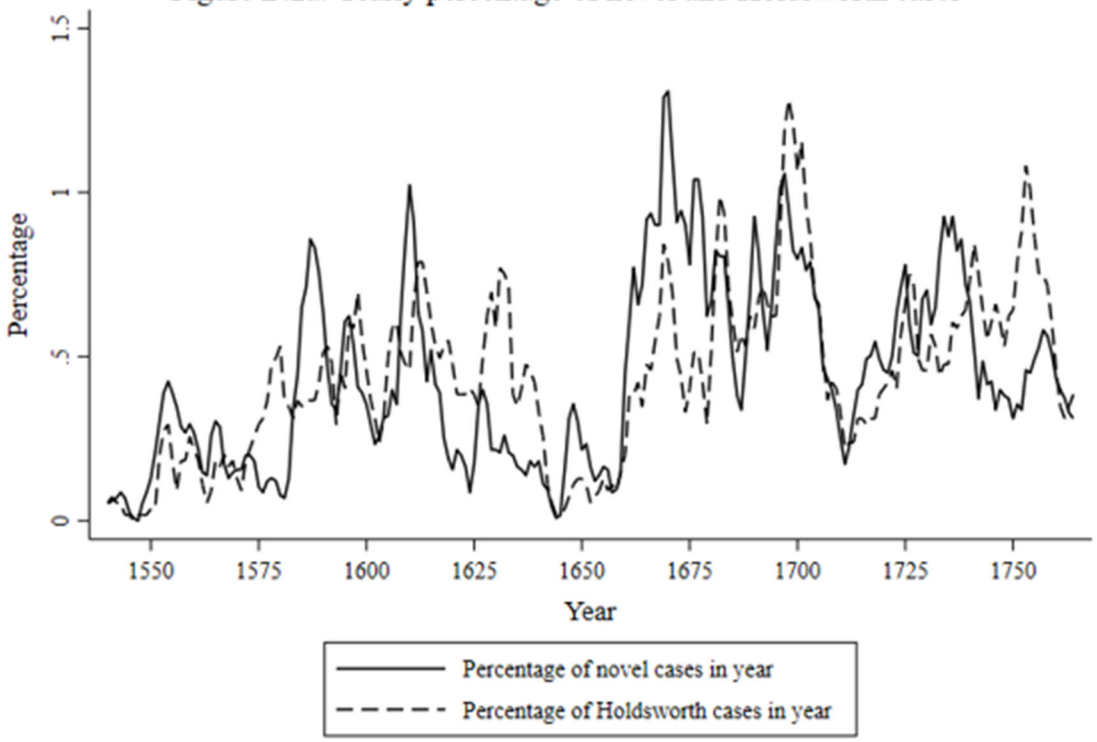

Figure B.2b: Cumulative sum of percentage of novel and Holdsworth cases, yearly

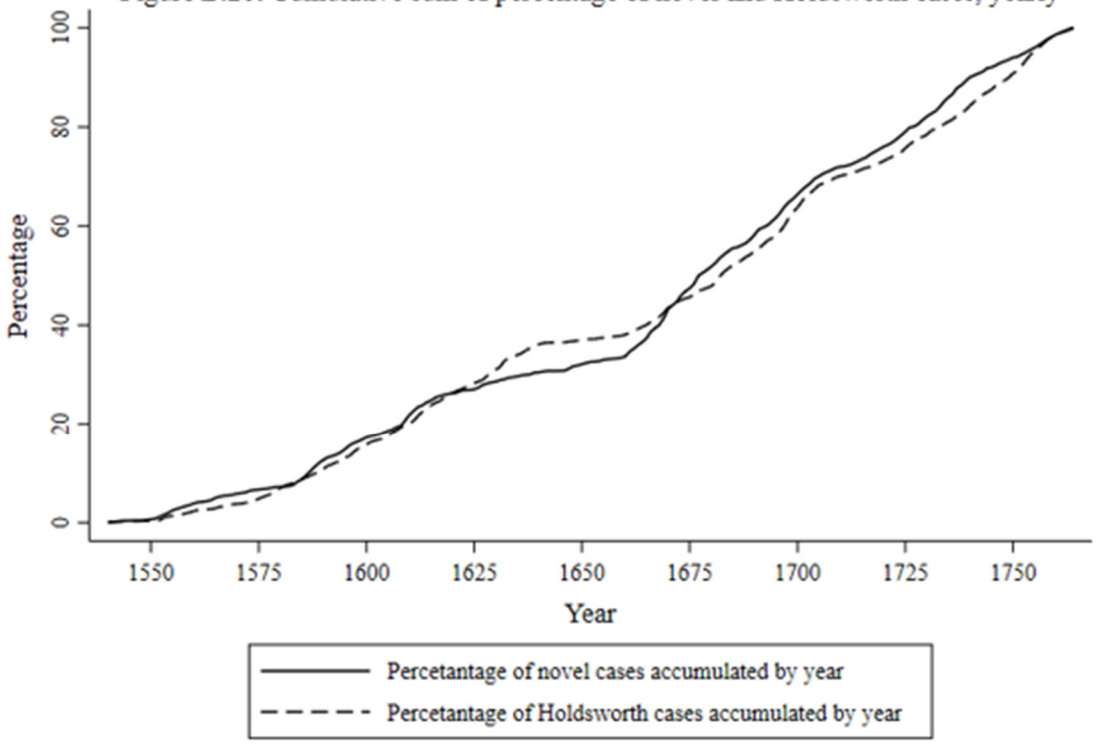

Figure B.3a: Yearly percentage of novel and Baker cases

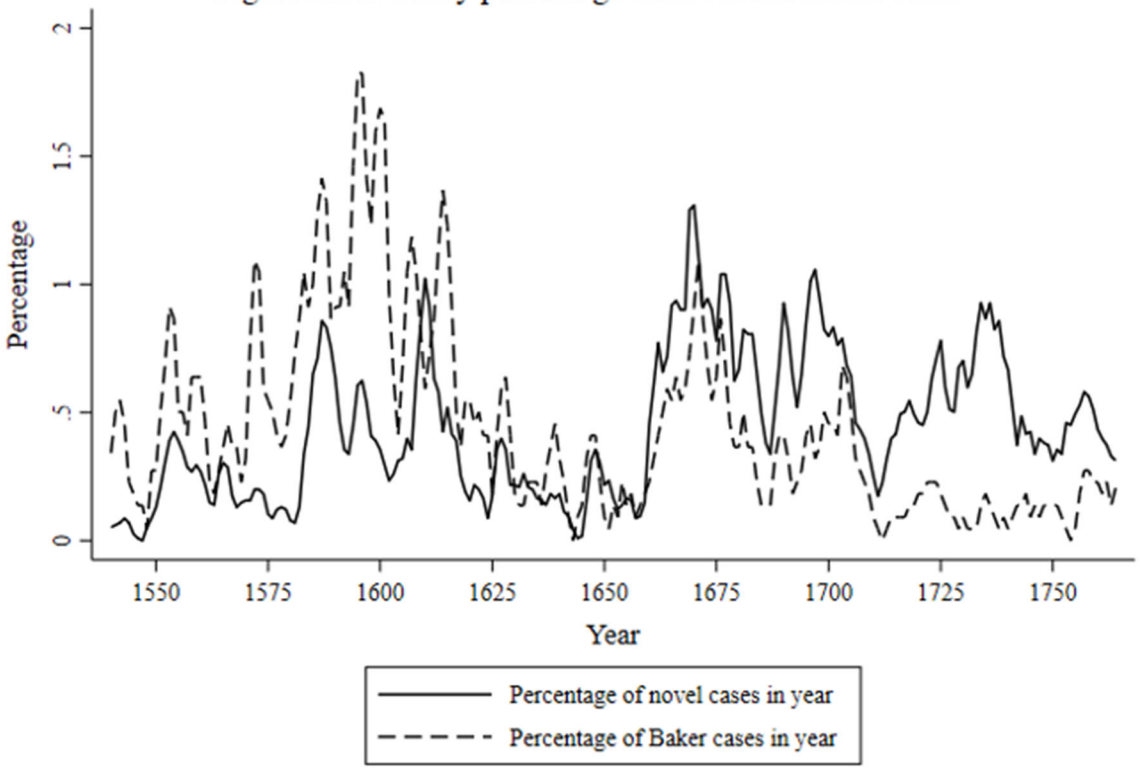

Figure B.3b: Cumulative sum of percentage of novel and Baker cases, yearly

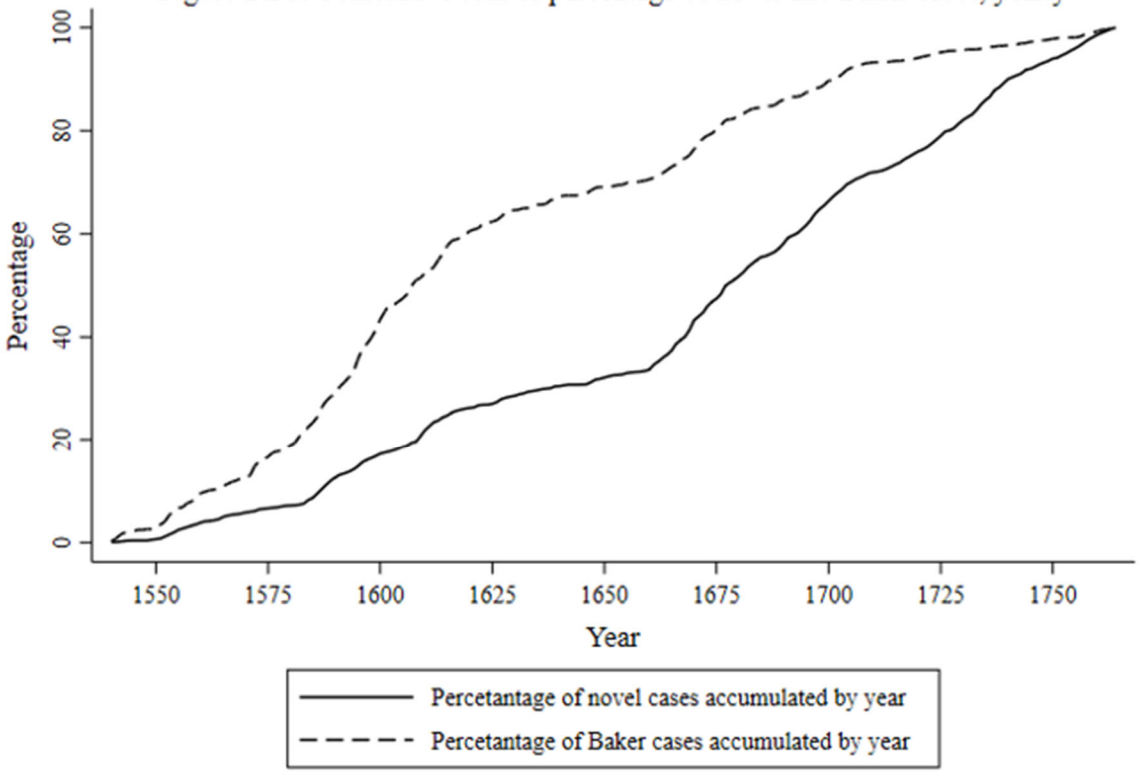

Appendix C: Growth in attention to precedent in the English Reports

Let p_{ijt} be the topic proportion for Precedent in document *i* produced by court *j* at time *t*. (All the relevant data can be obtained following the instructions in Appendix A.) Then attention to Precedent in court j in year t is:

$$a_{jt} = \frac{1}{n_{jt}} \sum_i p_{ijt}, \tag{C.1}$$

where n_{jt} is the number of case reports originating in court j in year t and the summation is over all case reports i originating with court j in year t. The time series of a_{jt} appear in the following figure.

Figure C.1 Attention to precedent by court

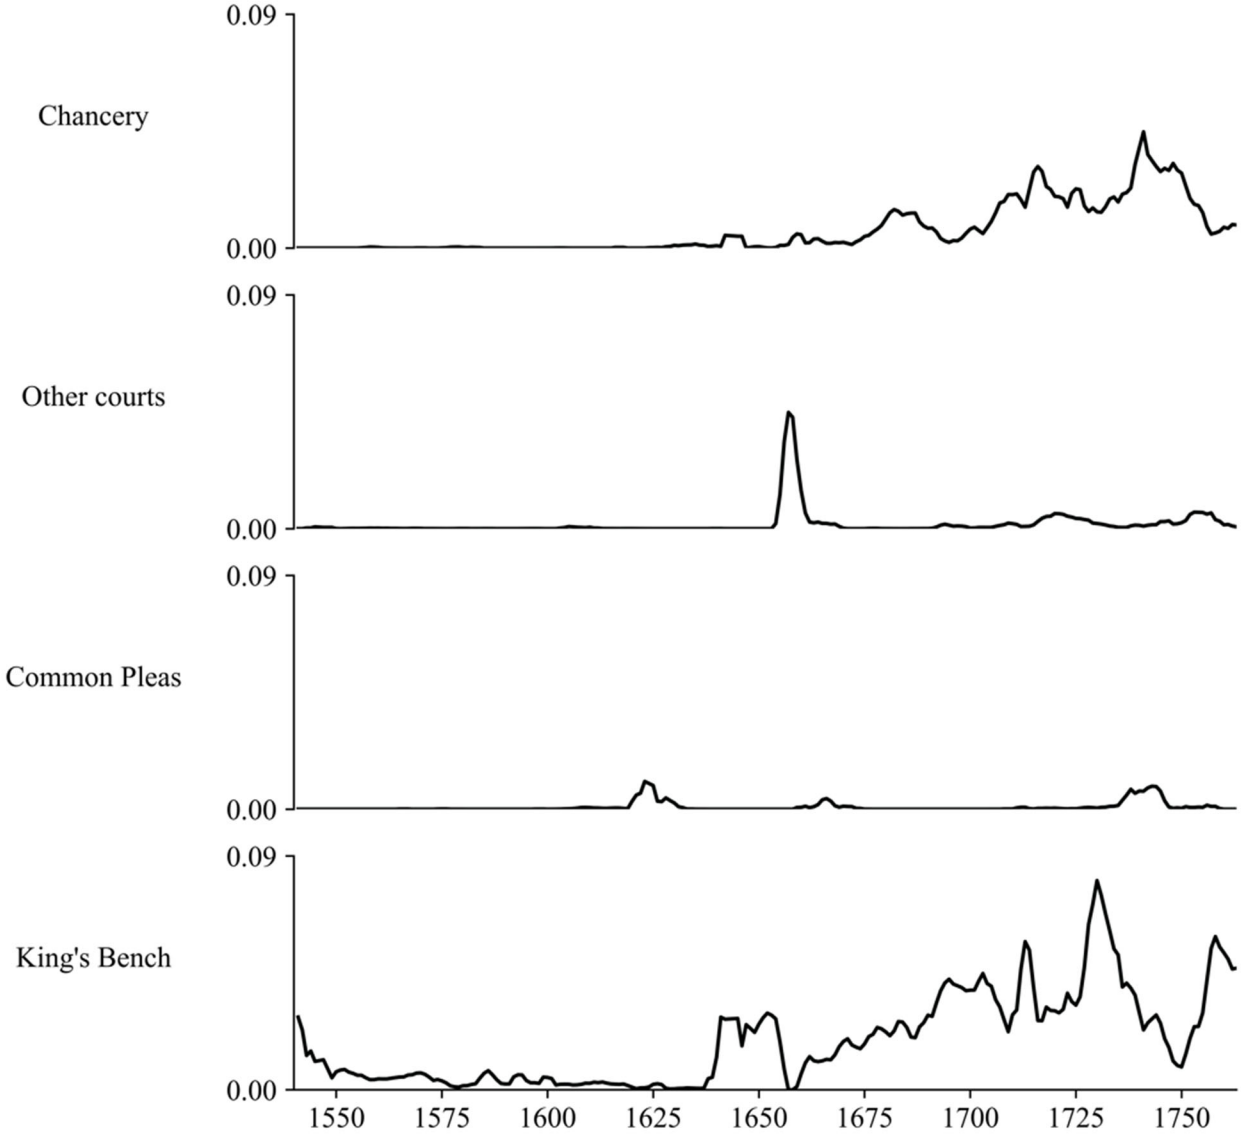

Appendix D: The legal themes

Table D.1: The allocation of topics to legal themes

Contract and Debt
- Assumpsit
- Bonds
- Identifying Contractual Breach
- Employment of Apprentices & Servants
- Executable Purchase Agreements
- Contract Interpretation & Validity
- Length & Expiry of Leases
- Rental Payments
- Repaying Debt
- Bankruptcy
- Prioritizing Claims
- Claims from Financial Instruments
- Pleadings on Debt
- Mortgages

Families and Inheritance
- Marriage Settlement
- Minors & Guardians
- Daughters' Legacies
- Geographic Settlement of Children
- Rights of Married Women
- Disentangling Heirs
- Specifying Inherited Property Rights
- Implementing Ambiguous Wills
- Contingency in Wills
- Execution & Administration of Estates
- Intestacy
- Validity of Wills
- Excluding Beneficiaries of Wills
- Estate Tail

Jurisdiction
- Inferior-Court Jurisdiction
- Equitable Relief
- Geographic Jurisdiction of Laws
- Equity Jurisdiction
- Prohibiting Jurisdiction

Markets and organizations
- Publishing & Copyright
- Regulating Commerce
- Municipal Charters
- Governance of Private Organizations
- Restraints on Trade
- Negotiable Bills and Notes

Politics
- Local Administrative Appointments
- Dignitaries
- Rights of Public Office
- Royal Patents & Tenures

Procedure
- Reviewing Local Orders
- Rendering Judgement
- Equity Appeals
- Arbitration & Umpires
- Interacting in Court
- Procedural Rulings on Actions
- Mistakes in Court Records
- Procedural Bills
- Writs of Error
- Jury Procedures & Trials
- Motions
- Court Petitions
- Correct Pleas
- Procedural Rulings on Writs
- Rulings on the Calendar
- Evidence Gathering & Admissibility

Sources of law
- Precedent
- Statute Applicability
- Clarifying Legislative Acts
- Contrasting Cases & Statutes

Real and personal property
- Timing of Property Rights
- Competing Land Claims
- Elizabethan Land Cases
- Equitable Waste
- Manorial Tenures
- Possession & Title
- Self-Help in Real-Property Disputes
- Common-Land Disputes
- Shared & Divided Property Rights
- Transfer of Ownership Rights
- Conveyancing by Fine
- Tree Law
- Uses
- Implementing Trusts
- Bailment
- Ownership of War Bounty
- Trespass to Goods

Other
- Revocation
- Determining Damages & Costs
- Multiparty Cases
- Attorney- & Solicitor-General
- Nuisance
- Actionable Defamation
- Wrongful Possession
- Indicting for Murder
- Habeas Corpus
- Decisions After Conviction
- Ecclesiastical Appointments
- Temporal & Spiritual Jurisdiction
- Tithes

Note: The names of themes appear in italics and the topics included in the themes are listed immediately below those names.

Figure E.1: Effect of time between cited and citing case on relationship between novelty and citations by novelty quartile