

RESEARCH ARTICLE

A machine-learning history of English caselaw and legal ideas prior to the Industrial Revolution I: generating and interpreting the estimates

Peter Grajzl^{1,2} and Peter Murrell^{3*}

¹Department of Economics, The Williams School of Commerce, Economics, and Politics, Washington and Lee University, Lexington, VA, USA, ²CESifo, Munich, Germany and ³Department of Economics, University of Maryland, College Park, MD 20742, USA

*Corresponding author. Email: pmurrell@umd.edu

(Received 21 April 2020; revised 29 June 2020; accepted 30 June 2020; first published online 5 August 2020)

Abstract

The history of England's institutions has long informed research on comparative economic development. Yet to date, there exists no quantitative evidence on a core aspect of England's institutional evolution, that embodied in the accumulated decisions of English courts. Focusing on the two centuries before the Industrial Revolution, we generate and analyze the first quantitative estimates of the development of English caselaw and its associated legal ideas. We achieve this in two companion papers. In this, the first of the pair, we build a comprehensive corpus of 52,949 reports of cases heard in England's high courts before 1765. Estimating a 100-topic structural topic model, we name and interpret all topics, each of which reflects a distinctive aspect of English legal thought. We produce time series of the estimated topic prevalences. To interpret the topic timelines, we develop a tractable model of the evolution of legal-cultural ideas and their prominence in case reports. In the companion paper, we will illustrate with multiple applications the usefulness of the large amount of new information generated by our approach.

Key words: Caselaw; English history; idea diffusion; institutional development; machine learning

JEL Codes: C8; N00; K10; Z10; P10

1. Introduction

The precocious socioeconomic development of England is standard fare both in books for broader audiences and in works aimed at professional economists, political scientists, and historians.¹ Many studies have linked the beginnings of that development to the creation of effective institutions. But amid those studies, one area stands out for its lack of detailed quantitative information: the development of English caselaw, and more generally the evolution of ideas embodied in the caselaw itself. This is remarkable given that caselaw is taken to be a defining characteristic of the English legal family. But it is also hardly surprising in view of the fact that caselaw develops via decisions in thousands and thousands of cases, which resist conventional quantitative analysis. Perhaps for this reason, research on the accumulation of caselaw and the associated legal ideas has not been the focus of those seeking evidence on the process of historical political and economic development, especially economists (Ogilvie and Carus, 2014).

Our objective is to generate and analyze the first-ever quantitative data on the development of English caselaw and the associated legal ideas from the mid-16th century to the mid-18th century.

¹See, for example, Voltaire (1734), Schama (2003), Fukuyama (2011), Acemoglu and Robinson (2012), North and Weingast (1989), Hayek (1960), Hoppit (2017), and Coffman *et al.* (2013).

We provide readers with time series capturing developments in 100 clearly defined areas of law. To accomplish this goal, we assemble a large-scale text corpus and utilize machine-learning tools for analysis of text-as-data. The integrated set of steps that we undertake provides a template for the way in which one can extract data from historical texts to generate a quantitative picture of micro-institutional change. The estimates that we generate should provide an essential input into any study that aims to investigate the history of England's institutional development.

Before describing in more detail what we do, it behooves us to explain how our work differs from the very detailed research undertaken for many centuries by legal historians, of which the most recent example is the masterful textbook of Baker (2019). The work of legal historians is based on conventional text analysis and, as a consequence, must rely on relatively small samples of cases. Therefore, that work pays most attention to landmark cases and important statutes. However, caselaw not only reflects such crucial developments, but also embodies thousands of more mundane cases, which fill in details. A case might, for example, decide that a certain type of financial instrument is, in principle, negotiable, thereby providing a watershed decision. But subsequently many other cases will focus on such mundane issues as whether only the bearer can receive payment; how indorsement must be done; which variants in the instrument are acceptable; which particular merchant practices are declared customary and therefore part of the common-law; whether the instrument must contain reference to those practices; and numerous other questions (Holden, 1955). A satisfactory body of caselaw exists only after such questions are answered by decisions in cases subsequent to the landmark ones, leading to myriad rules that together constitute pertinent caselaw. Then economic agents can comfortably use the new financial instrument because they know the exact configuration of the shadow of the law. Thus, the works of legal historians cannot substitute for the kind of analysis we pursue. These works provide qualitative histories that focus on crucial highly specific developments, whereas we provide a macroscopic quantitative overview that summarizes the flow of legal ideas through 52,949 cases decided by the high courts of England. Conversely, our work is no substitute for that of the legal historians. Our time series cannot tell us why particular turning points in the law happened and how the technical details of debate changed.

Sections 2 and 3 of this paper focus on the process by which we extract 100 time series of the development of English caselaw and legal ideas from 52,949 raw text documents. Section 2 begins with data and methods. The raw data are digitized versions of *The English Reports* (Renton, 1900–1932), which are generally regarded as the definitive set of case reports on English law from the mid-15th century to the mid-19th century. Because we are interested in developments before the Industrial Revolution, we focus on cases up to the mid-18th century (to be precise before 1765). The raw case reports are not suitable for statistical analysis because they contain numerous spelling variants, many early-modern English words that were later replaced by newer forms, and a significant amount of non-translated Latin. Therefore the first step in our analysis was to standardize the English orthography, convert old words to their modern versions, and translate the Latin, using Python programs and databases scraped from the web. The resultant case reports were then analyzed using the structural topic model (STM; Roberts *et al.*, 2014, 2016, 2019), a variant of the topic models that are now increasingly popular in the social sciences (Blei *et al.*, 2003). Section 2 of the paper closes by discussing the unique aspect of STM, which we use intensively in our paper, the integration of information (the metadata) about the documents themselves (i.e. case reports). This aspect of STM leads seamlessly to the production of the time series that track micro-institutional evolution.

STM estimates topics. As we discuss in Section 3, we find that our set of case reports is adequately summarized by 100 topics. Each of these topics captures the ideas in a clearly delineated area of substantive caselaw and corresponding legal thought. We assign topic names and interpret the nature of the topics using the words that are most associated with a topic and the documents that use the topics intensively. Examples of these topics are Habeas Corpus, Claims from Financial Instruments, and Tree Law. We provide examples of the process used to name all of the hundred topics: an appendix provides

details justifying all 100 topic names.² Therefore, the information contained in 52,949 case reports is aggregated into 100 clearly identifiable phenomena.

Given estimates of the 100 topics and a date attached to each reported case, one can ask how much of the corpus was devoted to each topic during any specific time-period. That is, one can estimate topic prevalence conditional on time. Thus, we generate 100 time series of topic prevalences. These constitute our central empirical contribution. An appendix summarizes the time series with 100 timelines – figures showing the prevalence of the use of each topic in the reports on cases heard in the high courts of England in each year from 1550 to 1764.

When perusing these timelines, a startling problem presents itself. The most common shape in the timelines is an inverted-U, with the tails showing topic prevalences quite close to zero. A tempting interpretation of these shapes is that a legal idea was first unknown, then entered the law and was much discussed, and then vanished from the law. This type of interpretation has been used in many contributions that examine changes in word frequencies. But this interpretation, we argue, cannot be correct. To show this point, we reflect upon the timeline and the legal history of *assumpsit*, a legal doctrine concerning which promises are implied in a contract. The historical literature persuasively argues that this fundamental idea became more and more accepted in the late-16th century and became authoritative early in the 17th century. In contrast, the topic prevalence for *Assumpsit* approaches zero by the early-18th century. We know therefore that the prevalence of a topic cannot possibly be monotonically related to the acceptance in English law of the set of ideas underlying the topic.

This leads us to generate a simple model of the way in which ideas are selected to appear in *The English Reports*, a model that should be broadly applicable to many document collections. In that model, we use an evolutionary replicator equation that captures the spread of ideas from new devotees to those initially skeptical, showing how the prevalence of the idea in case reports varies over time as the proportion of devotees and skeptics changes.³ This model can explain why an inverted-U topic-prevalence timeline will occur for a legal idea that is eventually completely accepted by the legal profession. Section 4, the final section of the present paper, develops the model and generates illustrative examples of its application, providing templates that facilitate interpretation of all the varied profiles that appear in the timelines for all 100 estimated topics.

In a companion paper (Grajzl and Murrell, 2021), we will build on the current analysis, to provide examples of the use of our data to elucidate the flow of history.

2. Data and methods

The dataset

The source of our corpus is a digitized database of *The English Reports* (Renton, 1900–1932), consisting of 129,042 reports of decisions rendered in the English courts of law between the early-13th century and the mid-19th century. Coverage in the early years is very incomplete, but by the mid-16th century the number of reports becomes substantial. *The English Reports* certainly do not cover all cases considered in the English courts, nor even all cases considered by the superior courts of the land. But one could argue that *The English Reports* provide a fairly complete record of those cases from the mid-16th century to the mid-19th century that came to be used by lawyers as the basis for legal precedent. In Appendix A, we provide a discussion of the history of *The English Reports*, particularly focusing on the mechanisms that led to the selection of cases that would appear in the, now definitive, edition of the reports (Renton, 1900–1932). In that appendix we conclude that, for a study of the development of the English law that was used by later lawyers, the editors of *The English Reports*

²The appendices are available at:

<http://www.econweb.umd.edu/%7Emurrell/articles/AppendicesMachineCaselawJOIE.pdf>

³This is essentially the same idea as when one is trying to model the prevalence of a disease in a population and that prevalence is captured in an inverted-U curve, whereas the proportion of the population that has had the disease is captured in an S-shaped curve (Henrich, 2001).

were, with some exaggeration, correct in viewing their publication as the ‘Complete Verbatim Re-issue of the Decisions of the English Courts prior to 1866’.⁴ Although this statement would be completely unsatisfactory for a legal historian who was interested in studying the overall configuration of legal activity at any given time and the process by which certain decisions and legal rules came to be made, it is a reasonable approximation for a study like the present one, where we aim to provide a macroscopic overview of the development of the English law produced by the courts and used by later courts.⁵

From this database, we first eliminated the reports (5.4%) that were available only in Law French.⁶ We processed the remaining reports in a series of stages. First, we standardized the chaotic orthography of the English that spanned many centuries. To this end, we converted all reports into modern English orthography, the only feasible standardization possible, which, in addition, aids in the understanding of historic legal texts for the 21st-century scholar. Second, we converted obsolete early-modern-English words into current versions. Third, we translated individual Latin words, which were occasionally left in the original texts by the editors who produced the English editions. After implementing these initial processing steps, we dropped 1% of reports that were judged to be outliers in terms of our success in matching all words in the reports to entries in modern English dictionaries. Exclusion of these reports ensured a high degree of orthographic standardization and success in word-by-word Latin-to-English translation, as well as minimizing the impact of any OCR errors. Appendix B summarizes the details of the construction of the initial corpus and its processing.

Next, we assigned to each report values of metadata variables. Given our interest in the evolution of legal ideas, our primary meta-variable of interest is the year of the reported case. (Further meta-variables are incorporated into the estimation as clarified below.) The case year was readily ascertained for 93.1% of the reports in the full database. For the remaining 6.9% of cases, most could be dated by visual inspection of electronic copies of the original volumes of *The English Reports*. For some reported cases the year was imputed. For the imputed year, we used the year of a case that was closest in pagination to the reported case in question. In less than 0.01% of the reported cases these methods were not successful. For this very small number of reported cases, we assigned the imputed year by randomly drawing a year from a time-interval determined by the years of the cases covered by the volume in which the report in question was contained.

In this research, we use a set of 52,949 reports on cases heard before 1765. We chose the year 1765 as the cut-off for a study of the early development of modern English law for two reasons. First, this is the point, very roughly speaking, that marks the onset of the Industrial Revolution. Our study of legal development therefore examines the caselaw that provided an input into the Industrial Revolution rather than the caselaw that was a product of the Industrial Revolution. Second, the character of the reports changes at approximately the same time. Reporting became less haphazard, more regularized, and in a style much closer to that of modern reports (Baker, 2019; Veeder, 1901; Winfield, 1925). Reporting also became much more voluminous. Because our objective is to study the features of the caselaw itself and, as far as possible, to produce results that do not reflect either incidental changes in the system of reporting itself or the amount of reporting, we focused our analysis on a time period before the sea-change in reporting occurred.

Figure 1 shows the distribution of the reports based on the timing of the underlying reported cases.⁷ The resulting corpus consists of 31,057,596 character strings. The average report length is 620 words,

⁴See the beginning of the Index of Cases in Volume 177 of Renton (1900–1932).

⁵Baker (1971a: 51) makes essentially this point: ‘It is natural that practising lawyers should rely chiefly on the printed reports, as being the most accessible... But it is clear that legal historians, who wish not merely to retrace the law as it actually developed but to understand how and why particular decisions were made, and what effect those decisions were believed to have at the time, must begin to tackle the shelves of manuscript reports which occupy so many of our major libraries’.

⁶Law French was an archaic language, used only by lawyers and hardly spoken by the beginning of the sixteenth century. By that time, it had become an artificial, not a natural, language. Its idiosyncrasies and bizarre combinations of antiquated 16th-century French, English, and Latin preclude any efforts to use Law French documents in the current study.

⁷The drop in the number of reports in the middle of the 17th century reflects the tumultuous Civil War period when the activity of the courts naturally decreased before a surge in activity that came with the return of relative stability.




Fig. 1. The distribution of the reports based on the timing of the reported cases.

with standard deviation equal to 1,510 words. After importing the resulting corpus into R, we further processed the text of the reports in order to prepare the dataset for analysis. Using R's `textProcessor` and `prepDocuments` functions, we converted all words to lower case, applied the Porter stemming algorithm, and removed standard English stop words (natural language words with little meaning, such as 'and', 'the', 'a', and 'an'), numbers, words with fewer than three characters, words included in only one document, and punctuation. After the resultant processing, the final corpus consisted of 5,928,494 letter-based strings (word tokens).

Methods

We use the data to estimate a topic model. Complementing conventional textual analysis, topic models are especially suitable for investigation of large textual corpora, when the primary goal of the analysis is to provide a macroscopic overview of the emphases in a corpus. Topic modeling has been embraced by both humanities scholars and social scientists (e.g. Grimmer and Stewart, 2013; Hansen et al., 2018; Mohr and Bogdanov, 2013). Recent research illustrates the use of topic modeling as a productive lens for the study of institutions and culture (Blaydes et al., 2018; Grajzl and Murrell, 2019, 2020). In topic modeling, the researcher postulates a model of the data-generating process and then uses the data to determine the most likely parameter values in the model. Documents are conceptualized as bags of words: word order is disregarded. An unsupervised machine-learning algorithm then leverages the co-occurrence of word-use across documents to identify 'topics' (Blei, 2012). Thus, despite viewing the words within a document as fully interchangeable, the resultant estimates still reflect, at least partially, the semantics of the underlying documents (Reich et al., 2015).

Analytically, the estimated topics are probability distributions over the corpus vocabulary, whereas documents (i.e. case reports) are mixtures of topics. Substantively, the topics capture ideas that are commonly used within the corpus; for example, in our study one such topic might be (and is) 'rights of married women'. The topics themselves, however, are solely a product of model estimation. The identification of the substance of each topic – its naming – is left to the researcher who examines the words most closely associated with the topic and undertakes a close reading of the documents that feature the topic most prominently. Thus, topic modeling is an exercise in discovery. It is therefore ideally suited to the most traditional of the historian's tasks, to characterize the ideas existing in a specific time and place: topic modeling estimates elements of a culture.

The procedural features of early common law render topic modeling an especially suitable methodological approach for analysis of *The English Reports*. During the period under consideration, common law used an extensive writ system, a set of procedural rules that strictly and rigidly conditioned the ways in which legal action could be brought before a court (Baker, 2019: Ch. 4). Once settled, the ‘forms of action’ associated with specific writs ensured not only uniform application of procedure, but also enabled the disputing parties, judges, and reporters to focus their attention on the substantive issues for which a legal remedy was being sought. Unless procedural considerations were a point of contention in a disputed case, a reporter’s account of the case could naturally focus on the substantive issues, and the associated legal remedies, conceptualizations, and ideas. This uniformity and focus on substance render topic modeling particularly applicable.

We estimate an STM (Roberts *et al.*, 2014, 2016, 2019).⁸ Appendix C provides a detailed exposition of the general structure of STM’s data-generating process. STM extends the workhorse latent Dirichlet allocation (LDA) model (Blei *et al.*, 2003) along two key dimensions. First, STM integrates document-level metadata, such as case year, directly into topic estimation. Second, STM allows topic prevalences to be correlated across documents even conditional on metadata. These features improve topic identification and interpretability, as well as enhance the accuracy and efficiency of the estimation of the effect of meta-variables on topical prevalence.⁹ Thus, in addition to report year, we further incorporate into topic estimation report-level metadata on reporter name, identity of the adjudicating court, *The English Reports* volume number, and the total citation count accrued by the mid-19th century. To maintain focus on the core question of the evolution of legal ideas, we do not investigate the effect of these additional meta-variables on topical prevalence.

3. The estimated topics

Choosing the number of topics

Estimation of an STM requires initially choosing the (fixed) number of topics. There exists no universally agreed-upon approach for this decision (Roberts *et al.*, 2014, 2019). Given the size of the corpus, we first estimated a series of STM models, varying the number of topics between 50 and 200. We then examined standard measures of goodness-of-fit such as held-out likelihood and size of residuals (Roberts *et al.*, 2019; Taddy, 2012; Wallach *et al.*, 2009). Close inspection revealed that the 100–topic model fit the data especially well. We then compared the 100–topic model with other models that had also fit the data well, using scores on average semantic coherence (a measure of the internal consistency of the topics) and exclusivity (a measure of the extent to which topics in the model are distinguishable from each other). None of the other models clearly dominated the 100–topic model on these two criteria. This conclusion was confirmed when we compared the 100–topic model to other models using our own intuitive judgment on the ease of interpreting topics and distinguishing between them. We therefore chose to base our findings on the 100–topic model. We verified that none of our substantive findings were sensitive to small variations in the number of topics.

Interpreting and naming the estimated topics

Upon estimating the 100–topic STM, we interpreted and named the estimated topics. Recall that topics are distributions over vocabulary and that documents are distributions over topics. To interpret and assign names to the estimated topics, we therefore proceeded as follows. For each estimated topic, we examined two sets of key words – Highest Probability and Score – that STM identified as most

⁸Law (2016), Grajzl and Murrell (2019, 2020), and Rice (2019) use STM to analyze legal documents. Legal texts have also been studied using other variants of topic modeling (see, e.g. Livermore *et al.*, 2017) and other computational approaches (see, e.g. Funk and Mullen, 2018; Livermore and Rockmore, 2019).

⁹See Roberts *et al.* (2014: online Appendix) for a detailed analysis of the advantages of STM over LDA.

highly associated with the topic.¹⁰ The resulting words, or rather their stems (e.g. ‘statut’), are listed in Appendix D. In addition, we studied the documents (reports) identified by STM as featuring a given topic most prominently. Our interpretation and naming of the estimated topics therefore also incorporated an element of conventional text analysis.

To illustrate the process of interpreting and naming the estimated topics, we provide three examples. (Appendix E provides a justification of the assigned names for each and every one of the 100 estimated topics.) In our first example, which STM labeled as topic 22, key words highly associated with the topic include ‘assumpsit’, ‘promis’, ‘indebitatus’, ‘consider’, ‘forbear’, ‘non’, ‘assum’, ‘indebt’, ‘pay’, and ‘debt’.¹¹ Accordingly, one top document elaborates on a case entailing ‘Error of a judgment given in Rochester in an assumpsit, where the plaintiff declared, that the defendant being indebted to him in fifteen pounds, in consideration the plaintiff would give time to him for the payment... Upon non assumpsit pleaded, and found for the plaintiff, and judgment accordingly, error was assigned, for that it was not shewn how the debt accrued; for it was said, that a general indebitatus was not sufficient. But it was resolved, that generally indebitatus assumpsit is not sufficient where it is the ground of the action; as to say, whereas he was indebted to him in such a sum, he promised to pay, there he ought to shew how he was indebted...’ (Austin versus Bewley, Croke Jac 548, 79 ER 470). In another top document, the courts rule that ‘Assumpsit for goods sold and delivered, must alledge a promise to pay their value; a special request of payment; and the day and place when and where such request was made’ (Osbaston versus Garton, Croke Eliz. 91, 78 ER 350). Similarly, ‘A declaration in assumpsit upon an implied promise is good, though “in considerations inde” [i.e. “in considerations thereof”] be omitted’ (Tyler versus Bendlowe, 2 Shower KB 180, 89 ER 876). This topic is clearly about the types of promises that the law assumes to be implied in a contract to exchange property, goods, money, or services. We therefore named this topic Assumpsit.¹² From this point on, we distinguish references to topics using an initial upper-case letter.

In a second example, topic 88, key words include ‘case’, ‘reason’, ‘though’, ‘law’, ‘opinion’, ‘therefor’, ‘think’, ‘object’, ‘question’, ‘differ’, ‘court’, and ‘cite’. In one top document, a disputing party argues that, with regard to the specific legal situation at hand, ‘the principal case, reported in two of the books could not be law’ (Rex versus Bryan, 1 Barnardiston KB 331, 94 ER 224). In response, ‘The Chief Justice said for his part, he thought, the case, which is reported in those two books, could not be law; but tho’ it might not be law in the general manner, as laid down, yet still it was a good authority at least for the present opinion’. In another top document, the judges debate whether to overturn a prior decision: ‘The Chief Justice said now, that he had considered farther of this case, and was fully satisfied of his former opinion. He said the only cases, that inclined the other way, are those in Roll’s Abr. 675, and Salk. 337, but as to the case in Roll’s upon the face of it there appears to have been a mistake in it, and there is a much more true and perfect state of it in Yelv. 117, and as to the case in Salk that appears to have been mistaken likewise.... Judge Lee said he had looked into his notes for the case of Thomson and Lion, and could not find that he had taken any notice of it. He said that these writs of error ought to be allowed in open court; and he did not see any reason for allowing the present one; accordingly the rule was made absolute’ (Harris versus Burley, 2 Barnardiston KB 431, 94 ER 599). In yet another top document, the Lord Chancellor stressed that it would be ‘much better to stick to the known general rules than to follow any one particular precedent which may be founded on reasons unknown to us; such a proceeding would confound all property; and then citing the case of Lady Lanesborough and Fox, as of the strongest authority to the case in point...’ (2 Equity Cases

¹⁰Highest probability words are the words that are most common for a given topic, but are also non-exclusive, in the sense that they may be associated with any number of topics. In contrast, Score words are ones that are highly associated with a particular topic relative to their association with other topics. Score reflects the log frequency of a word for a given topic relative to the average log frequency of the word across all topics.

¹¹The numerical labeling of topics by STM is an artifact of the estimation procedure and has no significance for substantive issues.

¹²See Baker (1971a, 1971b, 2019), Ibbetson (1984), Lieberman (1989: Ch. 6), and Lobban (1991: Ch. 9) for more details on assumpsit.

Abridged 674, 22 ER 566). The topic is clearly about how to clarify which caselaw is pertinent in a given situation and was therefore named Precedent.

In a third and final illustrative example, topic 17, key words include ‘word’, ‘action’, ‘speak’, ‘say’, ‘lie’, ‘innuendo’, ‘slander’, ‘libel’, ‘false’, ‘knave’, ‘thief’, ‘witch’, ‘whore’, ‘rascal’, and ‘pox’. In one top document, the court considered ‘Action for these words Thou art a rogue, and a thief. After verdict it was moved in arrest of judgment, that an action lies not for these words, for they are too general. But the Court held, that for the word thief it is maintainable, unless it be coupled with other words, which prove it to be no felony intended’ (Robins versus Franks, Croke Eliz. 857, 78 ER 1083). In another top document, ‘An action upon the case was brought for words, Thou art an healer of felonies, and adjudged maintainable, because in Devonshire where, &c. Healer signifies the same as hider or concealer. And the proverb there is, The healer is as bad as the stealer’ (Pridam versus Tucker, Noy 133, 74 ER 1096). Yet another top document revolves around a ‘Case for saying of the plaintiff, He has got the, pox of a yellow-haired wench’. In this dispute, the court decided: ‘it being here said, that he had got them of a yellow-haired wench, ‘tis actionable’; for wench in common speech is taken for a whore, and therefore it shall be intended the French-pox’ (Lymbe versus Hockly, 1 Levinz 205, 83 ER 370). The topic clearly concerns which uses of words could lead to actions in court. We name this topic Actionable Defamation.

It is important to emphasize that we could easily identify the idea or ideas underlying each and every one of the 100 topic, and that we were able to generate topic names that fit snugly within existing concepts in the legal, historical, and traditional text-analysis literature. Unsurprisingly, many of our topic names (e.g. Assumpsit, Motions, Bankruptcy, Uses, Bonds, Mortgages, and Marriage Settlement), resonate closely with specific legal concepts and instruments covered at length in textbooks on the history of English law. But our estimates also uncover, and differentiate among, nuances of ideas about particular legal concepts or instruments. The legal instrument of a will, for example, is featured prominently in the corpus. Accordingly, our estimates unearth a series of distinct topics depicting a range of legal ideas concerning the use of wills (Implementing Ambiguous Wills, Contingency in Wills, Validity of Wills, and Excluding Beneficiaries of Wills). Other types of topics encompass substantive issues that cut across many conventional areas of law (e.g. Revocation, Determining Damages & Costs), reflect ideas pertaining to prominent historical episodes (e.g. Elizabethan Land Cases), or characterize specific reporting styles (e.g. Coke Reporting). Thus, not all topics identified by our STM are narrowly tied to concrete traditional legal concepts or instruments. This observation underscores both the usefulness of topic modeling and the breadth of our findings. Finally, there are topics such as Non-Translated Latin and Multiparty Cases, which reflect specific features of reports that are arguably of no legal-historical interest. In some ways, these particular topics play a role similar to that of control variables included in regression analyses: by absorbing non-substantive elements of reporting style, they sharpen the estimation of the timelines for the substantive topics.

Table 1 ranks the topics based on their prevalences in the corpus. Perhaps unsurprisingly, Precedent is the most prevalent among all estimated topics in the corpus. Among the remaining fifteen most prevalent topics, nine are inherently procedural in character (Motions, Interacting in Court, Correct Pleas, Writs of Error, Procedural Rulings on Actions, Procedural Bills, Rulings on the Calendar, Rendering Judgment, and Procedural Rulings on Writs). This suggests that procedural issues were at the center of the development of early-modern common law, a finding we explore in greater detail in Grajzl and Murrell (2021).

The estimated topics over time: the topic timelines

We next turn to our core results, estimation of the topic prevalences over time. Methodologically, we leverage STM’s use of metavariables and estimate the relationship between the STM-generated topic prevalence and the year of the reported case for each of the 100 identified topics, modeled as a spline. This estimation process provides machine-learning-based estimates of the emergence and evolution of

Table 1. Topics ordered on the basis of expected prevalence in the corpus

%	Topic name (rank 1–50)	%	Topic name (rank 51–100)
3.81	Precedent	0.79	Municipal charters
3.63	Motions	0.75	Revocation
2.65	Interacting in court	0.75	Disentangling heirs
2.58	Correct pleas	0.75	Claims from financial instruments
2.57	Modern reporting	0.74	Manorial tenures
2.51	Keble reporting	0.74	Transfer of ownership rights
2.34	Writs of error	0.72	Timing of property rights
2.04	Procedural rulings on actions	0.72	Ecclesiastical appointments
1.86	Procedural bills	0.70	Tithes
1.83	Rulings on the calendar	0.67	Arbitration and umpires
1.82	Rendering judgment	0.66	Rental payments
1.76	Elizabethan land cases	0.64	Rights of public office
1.70	Procedural rulings on writs	0.64	Identifying contractual breach
1.60	Shared and divided property rights	0.63	Contingency in wills
1.59	Repaying debt	0.62	Intestacy
1.56	Jury procedures and trials	0.62	Implementing trusts
1.53	Habeas corpus	0.62	Local administrative appointments
1.51	Assumpsit	0.62	Uses
1.48	Decisions after conviction	0.62	Executable purchase agreements
1.46	Reviewing local orders	0.60	Competing land claims
1.37	Prohibiting jurisdiction	0.60	Marriage settlement
1.36	Actionable defamation	0.59	Negotiable bills and notes
1.36	Evidence gathering and admissibility	0.58	Estate tail
1.35	Equitable relief	0.58	Ownership of war bounty
1.29	Bonds	0.57	Bailment
1.20	Vesey footnotes	0.56	Contract interpretation and validity
1.13	Statute applicability	0.56	Regulating commerce
1.12	Execution and admin. of estates	0.56	Mortgages
1.10	Coke's procedural rulings	0.55	Daughters' legacies
1.09	Equity jurisdiction	0.54	Indicting for murder
1.09	Inferior-court jurisdiction	0.53	Multiparty cases
1.04	Specifying inherited property rights	0.51	Restraints on trade
0.99	Pleadings on debt	0.51	Vesey reporting
0.99	Clarifying legislative acts	0.49	Employ. of apprentices and servants
0.98	Non-translated Latin	0.49	Tree law
0.97	Coke reporting	0.48	Court petitions
0.96	Wrongful possession	0.48	Bankruptcy

(Continued)

Table 1. (Continued.)

%	Topic name (rank 1–50)	%	Topic name (rank 51–100)
0.96	Trespass to goods	0.48	Geographic jurisdiction of laws
0.94	Attorney- and solicitor-general	0.48	Minors and guardians
0.92	Possession and title	0.46	Dignitaries
0.90	Rights of married women	0.46	Conveyancing by fine
0.89	Nuisance	0.37	Publishing and copyright
0.89	Implementing ambiguous wills	0.35	Validity of wills
0.87	Royal patents and tenures	0.34	Temporal and spiritual jurisdiction
0.87	Determining damages and costs	0.33	Self-Help in real-property disputes
0.85	Geographic settlement of children	0.32	Equity appeals
0.83	Prioritizing claims	0.31	Equitable waste
0.83	Length and expiry of leases	0.31	Excluding beneficiaries of wills
0.82	Mistakes in court records	0.28	Governance of private organizations
0.82	Common-land disputes	0.09	Contrasting cases and statutes

Notes: % denotes the topic's expected proportion (in percentage terms) of the whole corpus.

emphases within English caselaw and its associated legal ideas from the late medieval to early modern period. The resultant information bears on institutional developments of fundamental importance, and it is provided, as never before, in a quantitative form.

Figures F1–F15 in Appendix F provide the 100 timelines of topic prevalences. Each timeline depicts the estimated mean topic proportion for all cases decided in each year, together with their associated 95% confidence intervals.¹³ Before we turn to the details of estimated topics, however, we must first address a problem that arises in the interpretation of the timelines. Resolution of this problem provides an essential building block of all our subsequent interpretation of the estimates.

4. Interpreting topic timelines

Topics are distributions over words: Figures F1–F15 in Appendix F ultimately reflect changing word use over time. Many past studies have used raw word frequencies to interpret the presence of certain ideas in a culture and it is tempting to think of adopting this approach in the current context. However, doing so immediately presents a problem.

Consider Figure 2, on Assumpsit. A word-frequency-based interpretation of that figure suggests that the importance of the idea of assumpsit rises steeply in the late-16th and early-17th centuries, peaks in the 1630's, and then dissipates by 1670.¹⁴ Direct historical evidence from the evolution of caselaw is inconsistent with this interpretation. The years surrounding 1600 were a time of intense debate on assumpsit and the core idea was accepted in 1602 (Baker, 1971b, 2019; Ibbetson, 1984; Lieberman, 1989: 128). As is typical in the development of caselaw, however, there were many subsequent cases that tested the use of the idea in related areas of law (e.g. Baker, 2019: Ch. 21; Lobban, 1991: Ch. 9). Cases involving assumpsit were still arising and being reported upon in the latter half

¹³To aid the reader in processing the information in this multiplicity of figures, we have organized the timelines into 15 broader themes, each of which focuses on a specific broad area of law or of legal thought. We discuss these themes in detail in Grajzl and Murrell (2021), where we present and comment on the aggregate trends in the development of major areas of law.

¹⁴In using this word frequency analogy, it is nevertheless important to recall that the STM does not rely solely on word frequencies when estimating topics; STM also leverages covariances in word use across documents and explicitly allows topic prevalence to vary with document-level metadata.




Fig. 2. The timeline for the topic Assumpsit.

of the 17th century. At the same time, the use of legal actions based on assumpsit became more and more common-place as the 17th century proceeded: assumpsit certainly did not lose relevance in the legal discourse. Rather, by the latter half of the century, the ideas and procedures related to assumpsit were fully integrated into English legal culture and were an important element of law. Indeed, attempts to extend the applicability of action in assumpsit into a variety of settings, including commercial transactions, use of bills of exchange, and employment relations, continued well into the 18th and even early-19th century (Lieberman, 1989: 127–130; Lobban, 1991: Ch. 4, 9).

In sum, the frequency of use of Assumpsit in the corpus, as captured in Figure 2, reflects the intensity of debate on the idea at a given point in time rather than the degree to which the idea is accepted by the legal profession by that point in time. (The figure shows the estimated mean topic proportion with 95% confidence bands.) This presents a challenge in interpreting our previous findings in terms of the spread and use of legal ideas. That challenge is taken up in the following subsections, by developing a model that can generate the types of time profiles of topic-use that we have observed in *The English Reports*. Notably, the model provides a method of interpreting all of the generated timelines. The conundrum identified above for Assumpsit, for example, vanishes once we use the model's lens for interpretation. We begin with an informal summary of the main features of the model and then proceed to the presentation of the formal model.

Model overview

We consider interactions among the set of legal professionals who are regularly involved in (civil and criminal) cases in the courts of law, both lawyers representing clients and judges adjudicating. We focus on these legal professionals rather than on plaintiffs and defendants because our interest is in the material covered by *The English Reports*, not in ascertaining who wins or loses or which cases end up in court. As detailed in Appendix A, the reported cases are usually those that contain contentious, undecided elements of law and ones that contain important, precedential decisions. These are the cases that are of interest to judges and practicing lawyers and that serve an important role in the education of lawyers. At any given point in time, the reported cases will therefore necessarily reflect prevailing economic, political, and other societal concerns, even if that link is not made explicit in our model.¹⁵

Our setting revolves around the interaction of legal professionals who naturally vary with regard to their attitude toward the use of specific, exogenously given legal ideas invoked in debates on particular cases. With respect to any specific idea, some professionals are devotees who envisage wide applicability of the idea to pertinent cases.¹⁶ Others are skeptics, comparatively less eager than devotees to invoke the idea when conceptualizing the pertinent legal issues or to extend the applicability of the idea to as broad a range of applications. The discourse between the devotees and the skeptics thus

¹⁵As Lobban (1991: 15) argues: 'Law has a relative autonomy; and those who suggest that the law is wholly or mostly shaped by social considerations outside the legal forum would do well to remember that, in practice, legal constructs formulated and imposed in the artificial forum of the court often determine and shape those social identities and forces said to control the law'.

¹⁶As Chief Justice of the King's Bench, Lord Mansfield, for example, was keen on extending the applicability of assumpsit 'to a large number of commercial transactions' (Lieberman, 1989: 129).

entails coordination costs that affect the payoffs of the legal professionals. The legal professionals, however, do change their behavior based on observed payoffs. In particular, when being a devotee to a legal idea yields a higher payoff than being a skeptic, initial skeptics are more likely to change and become devotees (and *vice versa*). The population-wide share of devotees (and skeptics) is thereby subject to evolutionary dynamics, and the resulting spread of legal ideas from devotees to those initially skeptical is analogous to the spread of disease in epidemiological settings. In particular, the temporal dynamics of the share of devotees gives rise to an S-shaped pattern of diffusion, with changes in the number of devotees of an idea characterized by an inverted-U shape.

Reporting on a legal idea occurs in response to the needs of the legal community. Congruent with the process of compilation of *The English Reports*, the demand for reporting on a legal idea in our model is driven especially by recent devotees to the idea. Thus, the extent of reporting on an idea at a given point in time – the theoretical analog of estimated topic prevalence – reflects changes in the share of adherents to the idea at that point in time. Of course, specific legal issues often entail the use of multiple legal ideas, some already present and others discovered unexpectedly over time. The reporting dynamics on such (multiple-idea) legal issues then reflects the aggregate dynamics of multiple associated legal ideas, tracing out patterns that can feature multiple local extrema. Importantly, our simulations of the reporting dynamics based on the resulting framework are able to replicate all core qualitative patterns in topic prevalences over time that appear in our STM estimates. Our model therefore captures the core features of the topic timelines generated by STM. In what follows, we develop these points formally.

Model setup

The set of legal professionals in question is of constant size, but membership can change with deaths and replacements. The professionals differ in their stance toward the use of a specific legal idea, denoted by i . For any given idea i , a devotee (D^i) freely embraces the idea. In contrast, a skeptic (S^i) is hesitant about using idea i , preferring a more traditional framework. Any professional's type with respect to a particular legal idea i (D^i or S^i) may or may not coincide with that professional's type with respect to another idea $j \neq i$ (D^j or S^j).

Professionals in the population interact with each other in and out of court. Out-of-court encounters take place, for example, in debates in the Inns of Court. We assume pairwise random interactions. Then in the framework of evolutionary game theory (Nowak, 2006) the consequences of an interaction at any time give rise to the following payoff matrix pertinent to idea i .

	D^i	S^i
D^i	B^i	$B^i - C^i$
S^i	0	0

Thus, for example, a devotee receives a payoff of $B^i - C^i$ when interacting with a skeptic on legal idea i . We assume that mutual adoption of a specific legal idea gives rise to benefits $B^i > 0$ that exceed those (normalized to zero) obtained when both interacting parties are skeptics who reject the legal idea. This captures the notion that the pertinent legal idea emerged as a result of a need to address a pressing legal concern and, thus, that the old ways used in dealing with the concern are less satisfactory for the devotees. For example, in the case of assumpsit, the old ways did not allow the pursuit of claims by a route that forced the defendant to face a jury trial (Ibbetson, 1984: 295; Lieberman, 1989: 127–128). Similarly, with regard to trespass, the applicability of the corresponding legal action was gradually extended to a multitude of situations involving injury to a person or property, to settings where prior remedies had been deemed insufficient (Lobban, 1991: 52).

Interaction between devotees and skeptics involves non-trivial coordination costs that arise because of differing conceptualizations of the problem at hand. When the legal idea under

consideration is relatively new, such costs will normally be larger for the devotees than for the skeptics; thus we assume $C^i > B^i$. Devotees have the weight of tradition to fight, as was certainly the case in *assumpsit* (Baker, 2019: 216). Moreover, in one specific application of an assumption that is standard in institutional economics (North, 1990), law is a complex web of interconnected precedential rulings and procedures, into which the old ideas fit snugly, whereas new ideas are dissonant with traditional procedures. For example, *assumpsit* opened up new avenues to sue for debt and was seen by skeptics as depriving defendants of traditional forms of defense (Baker, 1971a, 1971b). Similarly, action on the case, one of the common-law forms of action that eventually saw the widest use, was initially met with ‘great doubt’ (Lobban, 1991: 52). Because the skeptics are seen as relying on traditional legal forms no matter with whom they interact, the payoffs of skeptics are normalized at zero for all their interactions.

Evolutionary dynamics of a legal idea

For any idea i , the proportion of devotees at time t equals q_t^i . Then, the expected payoff for devotees from a random encounter at time t is $\pi_t^{D^i} = q_t^i B^i + (1 - q_t^i)(B^i - C^i)$ and for skeptics $\pi_t^{S^i} = 0$. In each interaction, professionals find out, with some probability, the difference in expected payoffs between themselves and the randomly matched individual. The probability of finding out the difference in expected payoffs is assumed to be proportional to the difference in payoffs, equal to $\beta^i |\pi_t^{D^i} - \pi_t^{S^i}|$.¹⁷ When professionals find out that their behavior results in a lower expected payoff than the behavior of their counterparts, they change their behavior. Thus, the probability of changing from being a skeptic to being a devotee is $\beta^i (\pi_t^{D^i} - \pi_t^{S^i})$ when finding out that $\pi_t^{D^i} > \pi_t^{S^i}$ and 0 otherwise. Analogous reasoning holds for the probability of changing from being a devotee to being a skeptic.

This mechanism leads to the following replicator equation for idea i (Nowak, 2006):

$$q_{t+1}^i = q_t^i + q_t^i(1 - q_t^i)\beta^i(\pi_t^{D^i} - \pi_t^{S^i}), \tag{1}$$

or in the case of the above payoff matrix:

$$q_{t+1}^i = q_t^i + q_t^i(1 - q_t^i)\beta^i(B^i - (1 - q_t^i)C^i). \tag{2}$$

The resulting process has three equilibria:

$$q^{i*} = 0, \quad q^{i**} = 1 - \left(\frac{B^i}{C^i}\right), \quad \text{and } q^{i***} = 1. \tag{3}$$

Of these, only q^{i*} and q^{i***} are stable. Their respective basins of attraction are $[0, q^{i**})$ and $(q^{i**}, 1]$. The dynamics of q_t^i then depends on the initial value, q_0^i . If $q_0^i > q^{i**}$ and q_0^i is sufficiently close to q^{i**} , then q_t^i approaches the stable equilibrium q^{i***} by tracing out an S-shaped curve that is standard in the cultural diffusion literature (Henrich, 2001).¹⁸ If one interpreted Figure 2 as capturing $\Delta q_t^i \equiv q_t^i - q_{t-1}^i$, then this type of diffusion curve would exactly correspond to the relationship obtained between t and q_t^i (i.e. the integral of the curve over $[0, t]$).

¹⁷The scaling parameter ensures that the probability is in $[0, 1]$. A sufficient condition is $\beta^i < 1/B^i$.

¹⁸We implicitly assume that the initial share of devotees to any given idea is always just large enough so that q^i converges to the stable equilibrium featuring devotees only (q^{i***}) rather than to the stable equilibrium featuring skeptics only (q^{i**}). Even when the initial share of devotees is small, this assumption will be satisfied if a devotee’s payoff when interacting with a skeptic is not large in absolute value.

However, Figure 2 actually captures the prevalence of the topic in the reports on cases decided at time t . The next modeling step, therefore, is to provide the logic of the link between the topic prevalence timelines and Δq_t^i . The key to that link lies in the incentives of case reporters.

The dynamics of reporting on a legal idea

Appendix A on the production of *The English Reports* (Renton, 1900–1932) provides the background to our characterization of the reporting process. The reporting on court cases was the domain of a set of professionals who were either within the core group of legal professionals modeled above or associated with it, for example, professional reporters trained in the law together with their editors and publishers. Historically, many of the case reports were originally written by lawyers directly involved in court proceedings, particularly judges. However, the volumes of these reports often appeared in print on the initiative of editors and publishers who were not the original authors, but were nevertheless very much responsive to the demands of the legal community.

Those initiating the process of publishing the reports had two different sets of motives. First, there were those like Coke who wanted to provide documentation of the most important new cases, and especially their own interpretation of these cases. Success in either goal would only arise if later lawyers used these works. Second, there were the editors and publishers who were interested in obtaining profits and prestige, but could only do so if their publications were viewed by the legal profession as providing important documentation that could be used for making arguments based on precedent. Hence, there was a market for reports and the demand in that market was derived from the thirst of the legal profession for knowledge about new legal developments.

Separate the devotees of idea i at time t into recently converted devotees, i.e. professionals who became devotees between period $t-1$ and t , and non-recent devotees, i.e. professionals who were devotees of idea i already in period $t-1$. The proportion of new devotees in the profession is Δq_t^i . These new devotees would have been skeptics in $t-1$ and then encountered an existing devotee and recognized the difference in payoffs, thereby converting a $t-1$ skeptic into a t devotee. That new devotee would then want information that provided sustenance for the new point-of-view. This demand for information would then drive the decisions of reporters, publishers, and editors: it would be a function of Δq_t^i . At the same time, those who have been devotees for a longer time would continue to demand material on the new learning. This type of demand would be a function of q_{t-1}^i .

Let r_t^i be the reporting activity on a given legal idea contained in cases decided at time t .¹⁹ Then a simple reduced-form behavioral expression that links r_t^i to the size of the two subgroups of devotees is

$$r_t^i = a^i + b^i q_{t-1}^i + c^i \Delta q_t^i, \quad (4)$$

where $a^i \geq 0$, $b^i \geq 0$, and $c^i > 0$. a^i reflects the overall interest of the legal profession in legal idea i that is independent of the prevalence of devotees and skeptics. b^i and c^i , respectively, capture the influence of new and other devotees on the production of reporting on idea i . New devotees exhibit a greater amount of zeal for the new idea and a greater need to educate themselves on the new ideas, since other devotees will already have accumulated much knowledge. Thus, $c^i > b^i$. This implies that the greatest levels of reporting on a given legal idea will occur during times of the most pronounced growth in the number of devotees, and not when the idea has already been absorbed by a large majority of professionals. In fact, once a large part of the profession has become devotees, the level of

¹⁹For brevity, we will often use the phrasing ‘reporting at time t ’ to mean ‘reporting on cases decided at time t ’. This elides differences between the time of reporting and the time of cases, since our topic prevalence timelines are based on the dates of cases and not the dates of reports. We do so without loss of generality. An extension of the model that would allow for a lag between the decisions on cases and reporting on the cases would yield similar qualitative insights. Moreover, the actual reports included in *The English Reports* were generated at the time when the case was heard and it was the inclusion of those reports into the editorially-compiled report volumes that naturally proceeded with a time lag.

reporting on the legal idea need not be significantly higher than it was during times when the idea had only a small number of devotees.

The conclusions, so far, are threefold. First, we have resolved the apparent inconsistency remarked upon at the beginning of this section. Assumpsit fades from view in the data because the idea is fully accepted within the legal profession. Second, the overall level of reporting on a given legal idea (or more generally word-use frequency) is an inadequate proxy for the extent to which an idea has been adopted in the discourse among legal professionals. Third, examination of Figures F1–F15 in Appendix D reveals in a straightforward way the degree of intensity of the development of a legal idea at any particular time. Our timelines provide a very simple answer to the question of when emphases on various areas of legal development occurred.

Illustrative timelines of legal development

Figure 2 on Assumpsit provides a particularly simple pattern, one featuring a single inverted-U. But among Figures F1–F15 in Appendix F there are much more complicated timelines, with several local maxima and minima, or even none at all. A simple extension of the above framework can easily explain how these configurations arise.

To extend the framework, note that our STM-estimated topics often capture more than a single legal idea. (Assumpsit is one topic that happens to be largely driven by one core idea.) Rather, many of the estimated topics encompass multiple legal ideas. Thus, the conceptual counterpart of an STM topic is a broader legal issue, which can include several related legal ideas.²⁰ Each of the legal ideas can arise quite separately in court cases that are relevant to the one general legal issue. For example, the overall legal issue might be the processing of an apprehended suspect of a crime and the relevant legal ideas ‘writ of habeas corpus’ ($i = 1$) and ‘remand on bail’ ($i = 2$).²¹ We therefore index the ideas with two superscripts: $k \in \{1, 2, \dots, K\}$ refers to one of the K topics (i.e. legal issues) and $i \in \{1, 2, \dots, n_k\}$ refers to a particular legal idea within that topic. Thus, r_t^{ki} is the number of reports on legal idea i related to the more general topic k . Denote by R_t^k the reporting on topic k appearing in cases decided at time t . R_t^k is the sum of reporting on the specific legal ideas relevant to k :

$$R_t^k = \sum_{i=1}^{n_k} r_t^{ki} \quad (5)$$

Expressions (2), (4), and (5) govern the dynamics of reporting on a given legal issue on cases decided at time t . When additionally normalized by the aggregate level of reporting on all legal issues at a given point in time, these expressions provide a parsimonious model of the aggregate dynamics of the prevalence of legal ideas implied by an STM-estimated topic. In what follows, we illustrate the resulting evolutionary dynamics with a series of simulations. The simulations show the relevance of the above conceptual framework in capturing and interpreting patterns in the evolution of STM-estimated topic prevalences.

We simulate the dynamics over 200 time periods, approximately the number of years covered by our topic-prevalence timelines. We focus on the role of the initial proportion of devotees of a specific idea, a key parameter that governs the dynamics of reporting on the idea. For the purposes of illustrating the core features of reporting dynamics, we treat the remaining model parameters as purely ancillary. The core simulation results are shown in Figure 3. In each panel of the figure, time is measured on the horizontal axis and the level of reporting on cases decided at time t is measured on the

²⁰In this section, we therefore view a ‘topic’ and a ‘legal issue’ as synonymous.

²¹This conceptualization best fits estimated topics that capture specific legal issues of direct legal-historical interest. But the basic modeling structure applies also to estimated topics that predominantly reflect the nature of a particular era (e.g. Elizabethan Land Cases) or even a reporting style (e.g. Keble Reporting), where the notion of an issue and idea are inherently less strictly legal in character.




Fig. 3. Simulated timelines. (a) One-idea issue, with relevant idea featuring many initial devotees. (b) One-idea issue, with relevant idea featuring few initial devotees. (c) One-idea issue, with relevant idea featuring some initial devotees. (d) Two-idea issue, with 1st (2nd) idea featuring many (some) initial devotees. (e) Two-idea issue, with 1st (2nd) idea featuring some (few) initial devotees. (f) Two-idea issue, with 2nd idea emerging late.

vertical axis. Because our aim is merely to highlight qualitative patterns, the scale on the vertical axis is irrelevant and thus suppressed.

As indicated by equation (5), reporting on a legal issue reflects the evolutionary dynamics of a set of legal ideas. Intuitively, the patterns present in the timeline of reporting on a legal issue will critically depend on the number of distinct legal ideas involved in discussing the legal issue. We first illustrate the dynamics of reporting on legal issues dominated by a single idea. Examples of the timelines fitting this pattern include *Assumpsit*, but also *Motions* (Figure F10a), *Bailment* (Figure F12), *Equitable Waste* (Figure F11a), and *Actionable Defamation* (Figure F13).

Recall that whenever the share of initial devotees falls within, but near the lower bound of, the basin of attraction of $q^{i***} = 1$, the evolution of the frequency of devotees of a particular idea traces out an S-shaped diffusion curve. Moreover, the demand for reporting on an idea is dominated by new devotees. Thus, the dynamics of reporting on a one-idea issue will vary with the frequency of initial devotees of that specific idea. Panels a–c of Figure 3 illustrate this point. When the initial share of devotees of the idea is large, the extent of reporting on the corresponding issue decreases over time, as in panel a. (For a corresponding example from the topic-prevalence timelines, see *Ecclesiastical Appointments* in Figure F4 in Appendix F.) When the initial share of devotees of the idea is small and the other parameters are such that changing from skeptic to devotee occurs less frequently, the amount of reporting steadily increases over time, as in panel b and also in the timeline for *Bankruptcy* in Figure F3. Finally, when the initial share of devotees is small and the parameters are such that changing from skeptic to devotee occurs easily, one obtains the inverted-U shape as in panel c and *Assumpsit*.²² All three panels a–c reflect the S-shaped pattern of evolution of the

²²In panel a of Figure 3, $B^1 = 0.9$, $C^1 = 0.1$, $\beta^1 = 0.03$, $a^1 = 0$, $b^1 = 1$, $c^1 = 100$, and $q_0^1 = 0.85$. In panel b, $B^1 = 0.9$, $C^1 = 0.1$, $\beta^1 = 0.1$, $a^1 = 0$, $b^1 = 1$, $c^1 = 10$, and $q_0^1 = 0.11$. In panel c, $B^1 = 0.9$, $C^1 = 0.1$, $\beta^1 = 0.1$, $a^1 = 0$, $b^1 = 0.2$, $c^1 = 100$, and $q_0^1 = 0.20$.

frequency of devotees. However, the change in the share of devotees as depicted in panels a and b is slow and these panels only reflect one slice of the S-shape: panel a shows the beginning and panel b the end. In contrast, panel c captures the complete pattern.

We next illustrate how our framework can be used to interpret the dynamics of reporting on legal issues that involve multiple legal ideas. For tractability, in each simulation we set the number of ideas equal to two, but this framework can be trivially extended to scenarios with $n > 2$. As we demonstrate below, the timelines depicting the dynamics of reporting on multiple-idea issues can involve multiple local maxima and minima, a pattern that never arises for single-idea issues. A first setting assumes that each of the two legal ideas has a strictly positive number of devotees at all points in time, with the frequency of initial devotees different for the two ideas. Under this setting, the local maxima and minima in the reporting timeline reflect differences across ideas in the timing and magnitude of changes in the frequency of recent devotees.

Panels d and e of Figure 3 illustrate two stylized patterns. Panel d depicts the reporting dynamics for a two-idea issue when the first idea initially has relatively many devotees, and the second idea relatively few. Then, the first inverted-U resembles that in panel c and is driven by changes in the rate of increase in the frequency of devotees to the first idea. The subsequent inverted-U-shaped pattern, again akin to that in panel c, reflects the dominance of reporting on the second idea in later periods. This pattern occurs in the timeline for Restraints on Trade in Figure F8. Panel e, in contrast, depicts the reporting dynamics for a two-idea issue where the frequency of initial devotees of the first idea is smaller than in panel d, whereas the second idea again features few devotees. In panel e, the reporting dynamics exhibits two successive inverted-U-shaped patterns, with the relative heights of the inverted-U's reversed, as in the timeline for Intestacy in Figure F14.

In a second setting, initially only one idea has a strictly positive number of devotees. This setting is intended to capture the notion that not all legal ideas have been under development since time immemorial. Rather, some legal ideas are introduced only once their usefulness becomes apparent in the context of changing socioeconomic circumstances (e.g. the building of canals).

Illustrating this setting, panel f of Figure 3 depicts the reporting dynamics when the initial frequency of devotees of the first idea is positive and the initial frequency for the second is zero. The second idea emerges later: there is a discrete jump in the frequency of devotees of the second idea from zero to a (small) positive number at $t = 100$ that puts q_{100}^i into the basin of attraction associated with the stable equilibrium in which all professionals become devotees, eventually.

The early reporting dynamics in this scenario are driven by reporting on the first idea, giving rise to the first inverted-U, analogous to that in panel c. In contrast, the later reporting dynamics, which give rise to another inverted-U pattern, are driven by the effects of the newly emerging second idea. The resulting reporting dynamics features two successive inverted-U-shaped patterns, qualitatively similar to that depicted in panel e.²³ Contingency in Wills in Figure F14 provides an example.

Note that all patterns in the topic-prevalence timelines in Figures F1–F15 in Appendix F can be viewed as comprising combinations of the six patterns in Figure 3. Thus the modeling framework presented in this section provides a method for interpreting the intensity of the development of ideas pertaining to specific legal issues (i.e. topics) as captured by the timelines produced by STM. Interpreted through the lens of our model, the shapes of the estimated timelines also indicate which legal issues are largely dominated by one versus multiple ideas.

More broadly, the modeling framework we have presented should also provide a cautionary note on the use of any analysis that relates word frequencies to the levels of any meta-variable. The basic lesson emphasized above applies as much to cross-sectional comparisons between countries as it does to within-country time series analysis: a lower prevalence level for a particular word could possibly

²³In panel d of Figure 3, $B^1 = B^2 = 0.9$, $C^1 = C^2 = 0.1$, $\beta^1 = \beta^2 = 0.1$, $a^1 = a^2 = 0$, $b^1 = 0.75$, $b^2 = 0.25$, $c^1 = 75$, $c^2 = 25$, $q_0^1 = 0.5$, and $q_0^2 = 0.21$. In panel e, $B^1 = B^2 = 0.9$, $C^1 = C^2 = 0.1$, $\beta^1 = \beta^2 = 0.075$, $a^1 = a^2 = 0$, $b^1 = 2$, $b^2 = 1$, $c^1 = 120$, $c^2 = 90$, $q_0^1 = 0.35$, and $q_0^2 = 0.15$. In panel f, $B^1 = 0.9$, $B^2 = 0.95$, $C^1 = C^2 = 0.1$, $\beta^1 = 0.075$, $\beta^2 = 0.31$, $a^1 = a^2 = 0$, $b^1 = 2$, $b^2 = 1$, $c^1 = 120$, $c^2 = 90$, $q_0^1 = 0.30$, $q_t^2 = 0$ for $t < 100$, and $q_{100}^2 = 0.075$.

indicate that the word signifies a lower acceptance of a certain idea. But, it might also indicate, as in the case of Assumpsit, that the idea is so fully accepted that authors do not feel the need to use the words associated with the idea.

Finally, it is important to note that, unlike in the example of Assumpsit, diminishing topic prevalence need not always correspond to acceptance of the set of ideas implied by the topic name. The reader should not forget that our assigned topic names are merely evocative labels, intended to resonate with the legal issues included in estimated topics; topic modeling certainly does not capture the doctrinal nuances that percolate through the legal debates. Thus, for some timelines, it is in fact the complement of the ideas captured by the assigned topic label that gradually gained acceptance. In the case of Royal Patents and Tenures, for example, the underlying idea that became accepted by the profession and society is that the Crown cannot indiscriminately grant rights. On the other hand, in the case of Ecclesiastical Appointments, diminishing topic prevalence probably captures the fact that widespread reception of pertinent ideas occurred hand-in-hand with decreasing relevance of the underlying discourse as society changed. Ascertaining the precise contextual reasons for decreasing topic prevalence for all pertinent timelines is, of course, an exercise that extends beyond the scope of the current research. Our estimates, however, render such an investigation feasible for the prospective users of our results. A careful scrutiny of individual timelines coupled with conventional legal-historical research promises to generate valuable insight into legal developments within many specific legal domains.

In Grajzl and Murrell (2021), we will build on the current analysis to provide multiple examples of the usefulness of the large amount of new information generated in this paper.

References

- Acemoglu, D and J. A. Robinson (2012), *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, New York, NY: Crown Publishers.
- Baker, J. H. (1971a), 'New Light on Slade's Case', *Cambridge Law Journal*, **29**(1): 51–67.
- Baker, J. H. (1971b), 'New Light on Slade's Case', *Cambridge Law Journal*, **29**(2): 213–236.
- Baker, J. H. (2019), *An Introduction to English Legal History*, 5th edition, Oxford, UK: Oxford University Press.
- Blaydes, L., J. Grimmer and A. McQueen (2018), 'Mirrors for Princes and Sultans: Advice on the Art of Governance in the Medieval Christian and Islamic Worlds', *Journal of Politics*, **80**(4): 1150–1167.
- Blei, D. M. (2012), 'Probabilistic Topic Models', *Communications of the ACM*, **55**(4): 77–84.
- Blei, D. M., A. Y. Ng and M. I. Jordan (2003), 'Latent Dirichlet Allocation', *Journal of Machine Learning Research*, **3**(4/5): 993–1022
- Coffman, D., A. Leonard and L. Neal (eds.) (2013). *Questioning Credible Commitment: Perspectives on the Rise of Financial Capitalism*, Cambridge, UK: Cambridge University Press.
- Fukuyama, F. (2011), *The Origins of Political Order: From Prehuman Times to the French Revolution*, New York: Farrar, Straus and Giroux.
- Funk, K. and L. A. Mullen (2018), 'The Spine of American Law: Digital Text Analysis and U.S. Legal Practice', *The American Historical Review*, **123**(1): 132–164.
- Grajzl, P. and P. Murrell (2019), 'Toward Understanding 17th Century English Culture: A Structural Topic Model of Francis Bacon's Ideas', *Journal of Comparative Economics*, **47**(1): 111–135.
- Grajzl, P. and P. Murrell (2020), 'Characterizing a Legal-Intellectual Culture: Bacon, Coke, and Seventeenth – Century England', *Cliometrica*, forthcoming. doi: <https://doi.org/10.1007/s11698-020-00202-5>.
- Grajzl, P. and P. Murrell (2021), 'A Machine-Learning History of English Caselaw and Legal Ideas Prior to the Industrial Revolution II: Applications', *Journal of Institutional Economics*, forthcoming.
- Grimmer, J. and B. M. Stewart (2013), 'Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts', *Political Analysis*, **21**(3): 267–297.
- Hansen, S., M. McMahon and A. Prat (2018), 'Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach', *Quarterly Journal of Economics*, **2**(1): 801–870.
- Hayek, F. A. (1960), *The Constitution of Liberty*, Chicago, IL: The University of Chicago Press.
- Henrich, J. (2001), 'Cultural Transmission and the Diffusion of Innovations: Adoption Dynamics Indicate That Biased Cultural Transmission Is the Predominate Force in Behavioral Change', *American Anthropologist*, **103**(4): 992–1013.
- Holden, J. M. (1955), *The History of Negotiable Instruments in English Law*, London: University of London.
- Hoppit, J. (2017), *Britain's Political Economies*, Cambridge, UK: Cambridge University Press.

- Ibbetson, D. (1984), 'Sixteenth Century Contract Law: Slade's Case in Context', *Oxford Journal of Legal Studies*, 4(3): 295–317.
- Law, D. S. (2016), 'Constitutional Archetypes', *Texas Law Review*, 95(2): 153–243.
- Lieberman, D. (1989), *The Province of Legislation Determined. Legal Theory in Eighteenth-Century Britain*, Cambridge, UK: Cambridge University Press.
- Livmore, M. A. and D. N. Rockmore (2019), *Law as Data: Computation, Text, and the Future of Legal Analysis*, Santa Fe, NM: The Santa Fe Institute Press.
- Livmore, M. A., A. B. Riddell and D. N. Rockmore (2017), 'The Supreme Court and the Judicial Genre', *Arizona Law Review*, 59(4): 837–901.
- Lobban, M. (1991), *The Common Law and English Jurisprudence, 1760–1850*, Oxford, UK: Oxford University Press.
- Mohr, J. W. and P. Bogdanov (2013), 'Introduction – Topic Models: What They Are and Why They Matter', *Poetics*, 41(6): 545–569.
- North, D. C. (1990), *Institutions, Institutional Change, and Economic Performance*, New York, NY: Cambridge University Press.
- North, D. C. and B. Weingast (1989), 'Constitutions and Commitment: The Evolution of Institutions Governing Public Choice in Seventeenth-Century England', *Journal of Economic History*, 49(4): 803–832.
- Nowak, M. A. (2006), *Evolutionary Dynamics: Exploring the Equations of Life*, Kindle Edition, Cambridge, MA: Harvard University Press.
- Ogilvie, S. and A. W. Carus (2014), 'Institutions and Economic Growth in Historical Perspective', in P. Aghion, and S.N. Durlauf (eds.), *Handbook of Economic Growth* (Vol. 2), Amsterdam, Elsevier, pp. 403–513.
- Reich, J., D. Tingley, J. Leder-Luis, M. E. Roberts and B. M. Stewart (2015), 'Computer-Assisted Reading and Discovery for Student-Generated Text in Massive Open Online Courses', *Journal of Learning Analytics*, 2(1): 156–184.
- Renton, A. W. (1900–1932), *The English Reports, Great Britain, Parliament, House of Lords*, Edinburgh, UK: W. Green & Sons.
- Rice, D. (2019), 'Measuring the Issue Content of Supreme Court Opinions', *Journal of Law and Courts*, 7(1): 107–127.
- Roberts, M. E., B. M. Stewart and E. M. Airoldi (2016), 'A Model of Text for Experimentation in the Social Sciences', *Journal of the American Statistical Association*, 111(515): 988–1003.
- Roberts, M. E., B. M. Stewart and D. Tingley (2019), 'stm: R Package for Structural Topic Models', *Journal of Statistical Software*, 91(2): 1–40.
- Roberts, M. E., B. M. Stewart, D. Tingley, C. Lucas, J. Leder-Luis, S. Kushner Gadarian, B. Albertson and D. G. Rand (2014), 'Structural Topic Models for Open-Ended Survey Responses', *American Journal of Political Science*, 58(4): 1064–1082.
- Schama, S. (2003), *A History of Britain: The British Wars, 1603–1776* (Vol. 2), New York, NY: Random House.
- Taddy, M. A. (2012), 'On Estimation and Selection for Topic Models', Proceedings of the 15th International Conference on Artificial Intelligence and Statistics, pp. 1184–1193.
- Veeder, V. V. (1901), 'The English Reports, 1292–1865', *Harvard Law Review*, 15(1): 1–25.
- Voltaire (1734), *Letters on England*, Kindle Edition. Edited by H. Morley. 1894 Edition by Cassell & Co. Salt Lake City, UT: Project Gutenberg.
- Wallach, H. M., I. Murray, R. Salakhutdinov and D. Mimno (2009), 'Evaluation Methods for Topic Models', ICML '09 Proceedings of the 26th Annual International Conference on Machine Learning, pp. 1105–1112.
- Winfield, P. H. (1925), *The Chief Sources of English Legal History*, Cambridge, MA: Harvard University Press.