

# Scaling Lower Court Decisions\*

CHRISTIAN ARNOLD<sup>†</sup>  
Cardiff University

BENJAMIN G. ENGST<sup>‡</sup>  
University of Mannheim

THOMAS GSCHWEND<sup>§</sup>  
University of Mannheim

May 31, 2019

## ABSTRACT

Legal scholars and political scientists studying courts and their decision-making have a common interest: They want to understand not only the patterns of the decisions judges make, but they also seek to analyze the policy implications of those decisions. To do so, scholars need to address an important measurement challenge: they have to determine both judges and their written decision's latent location in a latent case-space. While previous work successfully compares the locations of legal decisions to ideal points of the respective judges for the U.S. Supreme Court, strategies that are based on judges voting behavior obviously do not work in legal systems where such data is not available. In order to construct a common case-space without recurring to voting data we develop a new scaling model and apply it to lower court decisions in Germany. Locating such decisions in a common space is not only a mere academic exercise, but also has several important practical implications – the most prominent being *forum shopping* and *forum selling*.

*Paper prepared for the Conference on Conference on Data Science and Law, June 7-8, 2019.*

---

\*NOTE: Previous versions of this paper have been presented at the *Second Conference on Empirical Legal Studies in Europe*, Leuven, 2018. We thank Jens Frankenreiter, Nils Schaks and Philip Schroeder for helpful comments. We also thank Felicia Riethmüller for her insightful and instrumental research assistance. We acknowledge financial support for SFB 884 (project C4) at the University of Mannheim from the *Deutsche Forschungsgemeinschaft* (DFG) of Germany. All data and computer code necessary to replicate the results in this analysis will be made publicly available on our webpages on publication.

<sup>†</sup>Lecturer, School of Law and Politics, Cardiff University, Museum Avenue Cardiff, CF10 3AX. (arnoldc6@cardiff.ac.uk).

<sup>‡</sup>PostDoc, Sonderforschungsbereich 884, University of Mannheim, B6, 30-32, D-68131, Mannheim, Germany (engst@uni-mannheim.de).

<sup>§</sup>Professor, Department of Political Science, University of Mannheim, A5, 6, D-68131, Mannheim, Germany (gschwend@uni-mannheim.de).

# 1 Why Should We Scale Lower Court Decisions?

Legal scholars and political scientists studying courts and their decision-making have a common interest: They want to understand not only the patterns of the decisions judges make, but they also seek to analyze the policy implications of those decisions. Previous attempts to compare decision locations in a common space relative to ideal points of the respective judges have been developed for the U.S. Supreme Court (Clark and Lauderdale, 2010). Strategies that are based on voting data to identify preferences do obviously not work in systems where such data is not available. For instance, most constitutional courts rarely, or never, publish individual votes (Kelemen, 2013; Raffaelli, 2012). The same is true in particular for decisions of lower courts. How can we estimate and compare positions of various actors in a common space in situations where the individual votes are not published? Answers to this question are of great academic importance in order to explain the variance across cases, judges and courts in common law and civil law systems.

We develop a new scaling model to estimate and consequently compare locations of various judicial decisions in a common case-space. In contrast to previous research we lay-out an estimation strategy that does not rely on leveraging published individual votes. Our statistical model extracts more substantively meaningful information from citation patterns than previous approaches. Our model makes use of citations of various legal sources such as codes, statutes, laws, decrees, precedent or other previous decisions that describe the legal arguments judges make in order to locate written opinions into a common space. Doing this requires modeling the extracted citation data by a particular data generating process that describes judges' citation behavior when drafting a written decision. Specifically, our model estimates the frequency that a decision  $i$  is citing a particular legal source  $j$  assuming that the closer the decision  $i$  to a legal source  $j$  in a common space, the more likely are judges to cite this legal source in their decision when drafting a legal argument. In order to do this, we first pre-select comparable decisions that we assume to lie within the same case-space and show that they can be mapped into a common space. We then describe a strategy based on standard algorithms from information retrieval how to further automate case selection and enlarge the set of comparable decisions within the same case-space while still obtaining valid

estimates. We provide several applications for such common case-spaces leveraging the cited legal sources of lower court decisions in Germany.<sup>1</sup>

In addition to comparing decisions and courts within a meaningful common space, locating lower court decisions in a case-space has important implications to study the consequences of the well-known phenomena of *forum shopping* and *forum selling* (Klerman and Reilly, 2016; Bechtold, Frankenreiter and Klerman, 2019). Forum shopping lawyers will strategically choose the court that is most friendly towards the litigant within a particular case-space. Prominent areas where forum shopping occurs within the German system include press law or patent law (e.g., Dahlkamp and Schmid, 2014; van Lijnden, 2016; Bechtold, Frankenreiter and Klerman, 2019). If litigants shop for the most friendly forum for their case, judges might consequently be motivated to compete for them. Thus, they try to “sell” their forum by following certain procedures (e.g., oral hearings by default) that are more or less litigant friendly in order to attract more cases (Bechtold, Frankenreiter and Klerman, 2019). Moreover, the decision history of similar cases might increase or decrease the chances that the court attracts further cases. Since citizens expect impartial judicial arbitration, observable forum shopping and selling tendencies can have important repercussions for the public’s trust in and the legitimacy of courts. Our method helps to identify and quantify such tendencies.

## 2 How to Measure the Content of Legal Decisions?

Legal scholars and political scientists studying courts and judicial decision-making have been talking past one another for too long. They developed their respective field-specific approaches to study the same thing: both are interested in analyzing the output of court decisions. In order to overcome unproductive scholarly divisions, the doctrinal-politics approach (Clark and Lauderdale,

---

<sup>1</sup>Perhaps surprisingly, court rulings in Germany are not freely available. *Juris* GmbH is a publishing company that provides a database of legal documents and information on which we rely on in this paper. The *JURIS* data is very well suited for our endeavor, because it comprises a complete and already digitized corpus of written rulings of all available German court decisions. The annotation of the corpus with meta data makes this data even more valuable. In short, the extraordinary data quality allows us to develop state-of-the-art measurement models to tap into this so far unused data source.

2010; Lax, 2011) develops a common ground that brings together political science priorities and legal concerns. At its core, the doctrinal-politics approach adapts the standard spatial model that is known in political science and economics to represent actor's preferences and decisions outcomes. The resulting case-space model (Kornhauser, 1992a,b; Landa and Lax, 2008; Lax, 2011, 2012) allows judges to have different preferences over legal rules as well. Ideological differences between judicial actors, judges or collegial courts, might play out not directly but through expressing those differences in terms of legal principles and rules. Very similar to spatial models in political science this might lead to different partitions in acceptable and non-acceptable regions of the case-space. In order to describe and explain judicial decision-making we have to locate the case-facts as well as the preferences of the relevant actors as points in such a case-space.

So far one of the most significant limitations to apply this framework is the availability of respective measures to map judicial decisions in such a case-space. While it is possible to classify decision outcomes with dichotomous categories (e.g., liberal-conservative, guilty-innocent, affirm-reverse), this strategy fail to measure judicial decisions on a more fine grained level – ultimately a necessary step to locate the judicial decision in a case-space. Consequently, scholars use simplifying assumptions based on theories of intra-court bargaining to locate a court's decision in a common space. For instance, scholars equate the latent political position – and not the legal position – of a decision with the position of the respective median member of a particular coalition of justices (e.g. Brouard and Hönnige, 2017; Carrubba et al., 2012; Hönnige, 2007; Sternberg et al., 2015) assuming that the position of the median is known or “inherited” by other actors such a parties, senators or presidents with known positions (e.g. Epstein et al., 2007; Hönnige, 2009). Those strategies are fairly indirect, though. Important for the location of a written decision when evoking such assumptions is not so much *what* is written but *who* wrote it.

An alternative approach is to directly analyze the content of what is actually written in the decisions. Of course, interpreting and classifying the meaning of various parts of potentially a large number of decisions seems like the way to go. Law scholars started to embrace content analysis methods a while ago (Hall and Wright, 2008). However, when hand-coding potentially

a large number of written decisions not only time is of essence – the reliability of hand-coded classifications is a daunting issue as well. In short, scholars cannot have it both ways – an analysis of a large number of written decisions and a reliable procedure to do that. Thus, trade-offs have to be made and justified. When relying on automatic methods to classify and scale written decisions exploiting new developments in natural language processing and machine learning, the analyst is still “further away” from the actual meaning of the text as a manual coder but neither time nor reliability of the coding procedure becomes an issue.

In principle two different kinds of information from the written decisions can be automatically processed for locating them in a case-space. On the one hand, analysts can leverage the choice of words in which a decision is expressed. Previous approaches used supervised and unsupervised text-scaling methods (Evans et al., 2007; McGuire and Vanberg, 2005; Dyevre, 2019) to locate decisions or legal briefs in a common space. Given the particularity of the language used in such decisions more research has to be done to identify a case-space and estimate valid locations of decisions therein.

On the other hand, one can leverage the information that is being cited in such a decision as being diagnostic for its location. Judges writing a decision are free to cite particular legal sources to justify their judgement. Thus, the chosen selection of legal sources that are cited represent in some way the position of a decision. Accordingly, judges would cite different legal sources when they decide a cases differently, i.e. when the decision has a different position in the same case-space. Judges rely on various legal sources, such as statute or code but also on precedent to establish a legal argument or use evidence from other judgments to justify a certain outcome. These citations are important information to uncover the legal *relevance* of a judicial decision in a citation network (Fowler et al., 2007; Winkels, de Ruyter and Kroese, 2011).<sup>2</sup>

Citation patterns have also been used previously to show that citations are neither random nor independent of policy preferences of the authors. Frankenreiter (2017) exploits the institutional setting at the ECJ – unlike in other courts two opinions are drafted, one by the *Advocate*

---

<sup>2</sup>Whalen (2016) provides an overview over many different conceivable applications of how to use citation patterns for empirical legal research.

*General* and another one by the *Judge Rapporteur*. He finds that judges tend to cite decisions of judges appointed by Member State governments that have similar preferences regarding European integration. This empirical finding for the ECJ is also interesting in more general sense. The data generating process of citing precedent in case-law systems, while not necessarily binding, might be motivated by a similar logic than in common-law systems.

Finally, we are not the first to leverage citation patterns to map the latent *position* of written opinions into a case-space. [Clark and Lauderdale \(2010\)](#) analyze citation patterns of search and seizure cases and freedom of religion opinions written between 1953 and 2006 and use this data to locate US Supreme Court decisions meaningfully in a case-space. While [Clark and Lauderdale](#) rely on voting behavior of the judges to estimate valid positions, we are going to relax this assumption. In the next section we lay out our measurement strategy how to locate judicial decisions in a common space when individual votes are not available or not part of the published record.

### **3 Locating Decisions Using Citation Data**

Locating lower court decisions in a common space requires a new scaling model. Existing approaches which are well established in common law countries typically rely on judges voting behavior. In the German civil law system, however, these individual votes are not available, because it does not know concurrent and dissenting opinions to a majority decision. We therefore leverage the frequency of citations of various legal sources in lower court decisions – i.e., we leverage references to other decisions and the legal norms they cite – to scale decisions made by lower courts. Building on [Barberá \(2015\)](#) – who proposes a model to scale the ideology of politicians and their followers on Twitter – we adapt this approach to the legal context: Legal decisions cite legal sources that define a citation network – data which we leverage for scaling courts’ decisions within a particular case-space.

In what follows, we assume that judges are able to express their judicial preferences over a clearly defined legal topic through the variation and the frequency by which they cite particular le-

gal sources. Subsequently, we leverage those legal sources to estimate the location of each decision in a one-dimensional space. The estimated locations of decisions allow us to identify similarities and differences between the decisions in a latent space (similar to [Clark and Lauderdale, 2010](#)). The case-space model is “a variant of and a supplement to the policy-space model” ([Lax, 2012, 767](#)) and is designed to understand the substance of judicial policy-making with the help of case-characteristics ([Kornhauser, 1992a,b](#); [Landa and Lax, 2008](#); [Lax, 2011, 2012](#)). In other words, the estimates in a case-space represent judicial decisions and the differences and similarities between estimates reflect upon the judicial-policy preferences derived from the respective legal arguments in those decisions. In order to estimate a decision’s location in a case-space we will outline in the next section how judges can use legal sources to develop legal arguments.

### **3.1 How Do Judges Cite?**

Judges cite particular legal sources when drafting their decision. We chose to infer the location of written decisions in a case-space on the basis of citation patterns. “Similar decisions” – i.e. decisions that are in close proximity in a case-space – use a similar argument structure: Similar decisions cite similar legal sources and similar decisions will also recur to these legal sources with a similar frequency. In result similar decisions display a typical combination of legal sources necessary to substantiate a particular legal view. Thus, even when moving away from analyzing the raw text corpus, the cited legal sources alone still contain some information about which legal argument are advanced in the decision ([Clark and Lauderdale, 2012](#)).

How do judges choose the legal sources they cite in decision? We argue that there are three different mechanisms that drive what and how often a source gets cited in a judicial decision. First, the substance matter dictates a core set of legal sources – and any judge would agree to refer to this authoritative set of legal sources. Any judicial decision on this legal topic refers to these core legal sources in a similar way. As a result, any judicial decision in a certain domain displays a similar set and a similar frequency of these core legal sources. In essence, those core legal sources help to define the legal substance of what is at stake.

Second, there are legal sources that are idiosyncratically cited. In contrast to core legal sources, those so-called idiosyncratic legal sources appear in only one written decision. Judges may cite idiosyncratic legal sources to account for the particular characteristics of a case. Thus, in a set of comparable judicial decisions, citations to idiosyncratic legal sources appear in only one of the judicial decisions, but not in all others.

Finally, a judge's choice to cite various legal sources reflects tendencies in her rulings, i.e. a particular hard or a particular soft stance on a case. A judge is more likely to refer to those legal sources that are in line with her reasoning when drafting a decision. This is why: On the one hand, sources that are in line with the spirit of their own decision helps her making a legally sound argument which is why she is much more likely to quote these benign legal sources. On the other hand, a judge tends to refrain from citing legal sources that depart from her own opinion, because they generate cognitive dissonance and, eventually, cause extra effort when justifying her argument. She refers to dissenting legal sources – if at all – only in passing. Thus, a judge eager to present a cohesive legal argument more often cites legal sources that are in line with her argument than legal sources that are opposing or questioning her argument. Legal decisions from the same legal domain will therefore cluster according to the stance they take. Legal decisions on one side of a spectrum will resort to one particular set of legal sources. In contrast, judicial decisions on the other side of a spectrum will cite those sources that help the respective judge make the argument in the other way.

To which extent does each of the three data generating processes reveal information about preferences within the case-space? We begin with the core legal sources that define the substance of a decision. Any judge – irrespective of her own legal consideration – will cite this canonical set of legal sources and uses them in the same way as long as the cases are comparable in substance. When estimating the location of a decision, these core legal sources are therefore less informative for the eventual location of the decision in a case-space while it nevertheless facilitates the identification of a particular case-space.

Next, idiosyncratic legal sources neither reveal information about the location of a decision

nor the substantive identification of a particular case-space. Identifying preferences on the basis of citation networks requires the comparison of citations among several judicial decisions. Legal sources that are cited by one decision alone lack the frame of reference for comparison. We therefore exclude these references from the analysis.

Finally, information that helps to identify the location of judicial decisions stems from all other legal sources – that is from legal sources that are cited with different frequency by more than one decision. Decisions with similar legal reasoning refer to the same legal source in a similar way and therefore cause informative legal sources to cluster. Or, put differently, legal sources that can be used to substantiate one line of reasoning appear more often together in decisions that take this stance. In contrast, a judge who departs from a standard argumentation in her line of reasoning is most likely to bolster her argument with new legal sources - anything that somehow can substantiate the claims made in the decision that pushes this decision away from the well-known authoritative argument in land-mark decisions. Suppose one decision  $D_1$  cites legal sources  $A$  and  $B$  several times and another decision  $D_2$  cites legal source  $B$  and  $C$  several times. Now suppose we analyze a third decision  $D_3$  citing only legal source  $B$  several times. We would conclude that  $D_3$  should lie somewhere between the first two decisions if all three decisions belong to the same case-space. Thus, we assume that “similar decisions” cite much the same legal sources and explicitly model these co-occurrences to identify decision locations in a given case-space.<sup>3</sup>

## 3.2 Implications and Features

There are a number of implications when estimating decision locations using citation networks. All estimated locations are relative to one another. Adding to or subtracting from the set of analyzed decisions under consideration may change the estimated locations already. Again, imagine a judicial decision  $D_1$  citing the legal source  $A$  and  $B$  and another decision  $D_2$  citing legal sources  $B$  and  $C$ . Now suppose we add a third decision  $D_3$  citing legal sources,  $B$ ,  $C$  and  $D$ . We would

---

<sup>3</sup>This is comparable to scaling models in the context of roll-call votes in legislative politics (e.g. [Clinton, Jackman and Rivers, 2004](#)). There, ideologically similar legislators vote ‘yes’ (or ‘no’) on much the same roll-call votes. Those who only sometimes vote ‘yes’ or ‘no’ are identified to lie somewhere in between those two groups of legislators.

conclude that the third decision  $D_3$  is more ‘extreme’ than the judicial decision  $D_2$ . What effectively happened is that  $A$  and  $B$  do not represent the minimum and maximum of the underlying case-space. Instead, the extreme locations of the new space is now defined by  $A$  and  $C$ . If we define the space to have a standardized metric – for example z-scores – then the numerical value from  $B$  would change to a more central position if citations to  $C$  are also analyzed.

Are all citations indicative of a legal preference for one side or another of a legal argument? Not every legal source is cited positively, i.e. that it is used to affirm an argument. Judges also develop counter-arguments in which legal sources are cited negatively in order to dispute another argument. Similar to the analysis of the use of precedent in decisions of the ECJ (Jacob, 2014), these negative citations are nevertheless used to bolster the credibility of the decision, demonstrate the argument’s coherence or to fend-off potential revisions from higher courts. Judges may distinguish their legal argument from existing ones in order to show why the existing ones are not binding or valid by entertaining a potential counterargument or introducing a more nuanced view (“ $D_1$  is similar to decision  $D_2$  but not to decision  $D_3$ ”). Positive and negative citations will not be mentioned with the same frequency, however. Given limited time and resources, judges are more likely to cite a legal source in line with her legal reasoning than a legal source that runs against it. Thus, judges will cite more legal sources in favor of her argument. And if they quote a legal source for a counter-argument, they will refer to this source only on a few occasions. In result, judges are not only more likely to cite legal sources in favor of their argument at all, judges will also refer to them more frequently.

In contrast to previous research (e.g. Clark and Lauderdale, 2010) we treat all citations alike and do not select particular legal sources, such as substantial but not procedural legal sources. It is not necessary to distinguish positive and negative citations and between different types of legal sources to estimate the location of lower court decisions in an appropriate case-space. Positive and negative citations are already accounted for with different expected frequencies in the data generating process. Negative citation should occur less often the positive citations. Furthermore, irrespective of the type of the legal sources, the frequency with which lower court judges use

certain legal sources, substantive or procedural, will always be informative if it reflects a choice of the judge. <sup>4</sup>

## 4 A Common Space Estimator for Decisions and Cited Legal Sources

In the previous section we developed a model under what conditions judges cite particular legal sources and how often. In this section we propose a statistical model that reflects our insights about the data generating process. Thus we need an appropriate probability model for the data we have and come-up with a systematic component of such a statistical model that reflects our theory about judges citation behavior.

The data we observe consists of a  $n \times m$  matrix  $\mathbf{Y}$  of citation counts, where element  $y_{ij}$  represents the frequency how often a judge writing decision  $i \in \{1, \dots, n\}$  is citing legal source  $j \in \{1, \dots, m\}$ . We use a Poisson distribution as a typical probability model for how often we expect decisions to cite particular legal sources. Furthermore, the systematic component of our model describing the expected number of citations  $\lambda_{ij}$  rests on two assumptions. (1) Each decision has a fixed location along a unidimensional case-space. (2) A written decision will cite a legal source more often the closer their locations are in this case-space. Given that we neither observe the locations of the decisions nor of the cited legal sources they have to be estimated in a unidimensional space. We use the following statistical model to model the probability of observing any particular distribution of citations to legal sources.

---

<sup>4</sup>We also do not distinguish between legal sources for pragmatic reasons. We explicitly want to show that it is possible to retrieve the positions of judicial decisions with minimal effort. Our goal is to use all available legal sources. Any further distinction between various types of data – even with the help of machine learning classifiers – would involve the hand-coding of at least a training set and thus increase the necessary effort for analysis.

$$y_{ij} = \text{Poisson}(\lambda_{ij}) \tag{1}$$

$$\lambda_{ij} = \exp(\alpha_j + \beta_i - \gamma \|\theta_i - \phi_j\|^2) \tag{2}$$

We use a quadratic loss function to operationalize both assumptions. The distance between a decision  $i$  and a legal source  $j$  is expressed as  $\|\theta_i - \phi_j\|^2$ , where  $\theta_i \in \mathbb{R}$  is the location of decision  $i$  and  $\phi_j \in \mathbb{R}$  is the location of legal source  $j$ . The estimated rate with which increasing discrepancies in locations between the decision and a legal source reduce the expected frequency of citations is captured by  $\gamma$ .

In contrast to [Clark and Lauderdale \(2010\)](#) we provide a richer parameterization of the systematic component of our model (similar to the model of [Barberá \(2015\)](#) in a different context) by adding the parameters  $\alpha_j$  and  $\beta_i$  to capture idiosyncrasies of the judicial decisions and the cited legal sources. The parameter  $\alpha_j$  expresses the authority of a legal source  $j$ . Some legal sources are by default cited more often than others, simply because they are more relevant. Similarly,  $\beta_i$  captures the decision-specific differences. Some decisions cite on average more legal sources than others.

Identification is a core concern when estimating latent variable models. These models typically have more parameters than observations which results in infinitely many ‘correct’ solutions. Shifting, rotating or scaling one parameter can be easily offset by inversely shifting, rotating or scaling another related parameter. A unique solution needs to address all three concerns. The model has to be anchored in the parameter space so that it does not float around. Identification also requires a well defined scale and the model must not be allowed to arbitrarily stretch. But even if local identification is guaranteed, it would still be possible to symmetrically rotate the model in the parameter space by inverting all parameters.

Bayesian MCMC estimators can address these identification challenges. It is possible to specify priors in a way that helps offset identification issues and provide near-identification ([Gel-](#)

man and Hill, 2007; Jackman, 2009). Following (Barberá, 2015), we identify our model locally with the following priors.

$$\alpha_j \sim N(0, 1) \tag{3}$$

$$\beta_i \sim N(\mu_\beta, \sigma_\beta) \tag{4}$$

$$\phi_j \sim N(\mu_\phi, \sigma_\phi) \tag{5}$$

$$\theta_i \sim N(0, 1) \tag{6}$$

A standard normal distribution defines  $\alpha_j$  and  $\theta_i$  and guarantees local identification. Global identification is a bit more challenging. The literature in political science can guide us here. Using item response theory (e.g., Clinton, Jackman and Rivers, 2004; Poole and Rosenthal, 2007) this literature estimates locations of political actors and legislative proposals in a common space. Rotational invariance is dealt with in at least two ways. Authors either globally identify their model. Highly informative priors on a well known political actor clearly determine who belongs to “the right” or “the left” (Clinton, Jackman and Rivers, 2004; Martin and Quinn, 2002) side of a unidimensional space. In a similar vein, in the case of multi-dimensional scaling where prior intuitions about political actors might not be as well defined, Jackman (2001) proposes to set priors on well understood legislative proposals instead. The other option is to not identify a model globally at all. The respective posterior distribution would have two modes. In theory, the sampler could visit both sides of the bi-modal posterior distribution. In practice, however, this typically turns out not to be an issue as long as there is a reasonably large number of political actors whose locations can be reasonably well (Jackman, 2001, 2004) distinguished. It is then sufficient to initialize the sampler around an educated guess. Even with symmetric – and thus for rotational invariance uninformative – priors the algorithm quickly converges to the correct posterior (Jackman, 2009).

As for the first strategy, fixing certain decisions or legal sources is not possible in our application. We propose our model in a context where the analyst is neither certain about a clear

location of certain decisions nor the location of the cited legal sources. Avoiding any unjustified bias from priors, the model should identify the parameters on the basis of the citation data only. We therefore generate starting values empirically from a pre-run of the model. In a first run, we draw 8 parallel chains 1000 times from the posterior distribution. Because each chain explores only one side of the bi-modal posterior, the chains have to be collapsed to one side. Whenever the median of a chain has a different sign than the median from the first chain, we multiply all draws for this parameter with -1 to map the posterior on the same side. Averaging over all 8000 posterior draws from this first run, we obtain a good proxy for a starting value for the second run. In this second run, we then draw 10'000 times from the posterior to effectively explore its central tendency and variation.

## 5 Decisions of Lower Courts in Germany

We showcase the model for a previously unexplored case – decisions by German *Landgerichte* (district courts) and *Oberlandesgerichte* (regional courts of appeal). The *JURIS* data base is very well suited for our endeavor, because it comprises a complete collection of all available German court decisions. Each data base entry contains not only the fully digitized text, but also further meta data such as titles, dates or information about the respective court; it provides information on the further fate of a legal proceeding in front of other courts, all citations a judicial decision makes (backward citation), as well as other judicial decisions that refer to a certain decision (forward citation). The annotation of the corpus of decisions with meta data makes the information in this data base valuable for our purpose. It allows us to easily identify and then extract the necessary citation data.<sup>5</sup>

We proceed in three steps. First, we predefine a set of seed decisions that belong to a uni-dimensional case-space. Estimates of the decision locations are only valid if the decisions are

---

<sup>5</sup>After closer inspection we found certain citations using the ‘Ibid.’ citation style (‘a.a.O’ in the German context) are not annotated. We are not the first to realize this (Coupette and Fleckner, 2018). Thus, the counts of existing links between a decision and a legal source is potentially lower if a judge uses this citation style. Nevertheless, this should make our estimates more conservative.

distinctly identifiable on this latent space. Second, using the meta data, we then construct the respective decision-source matrix  $Y_{ij}$  that identifies all cited legal sources ( $j$ ) for each pre-selected decision ( $i$ ).<sup>6</sup> Third, we show how to automatically enlarge our sample of seed decisions and validly estimate the locations of the larger sample of decisions and legal sources. In addition to that we also use present a detailed case study to show that two decisions with estimated similar locations provide similar arguments and, conversely, two decisions with estimated dissimilar locations employ different arguments. This further supports the validity of our estimates and the procedure we propose.

We decide to measure and identify a legal source using its section and paragraph – *Paragraph* and *Absatz* in the German context. A reference to the same law and section, but with a different paragraph would be counted as another legal source. Legal sources can be references to German Civil Code (e.g., a particular section of the BGB), to criminal law (e.g., a section in the StGB), to code of civil procedure (e.g., a section in the ZPO) or to a section in the German constitution (the *Grundgesetz*, GG). We consider previous decisions of both lower courts (e.g. a decision written at the OLG Hamburg) and higher courts such as the Federal Court of Justice (BGH) or the German Federal Constitutional Court (BVerfG) as a specialized court of higher order.

## 5.1 Application to Press Law

Our model identifies systematic differences in courts’ legal arguments. We chose to exemplify our model for the German case, because it is a particularly challenging one. Typically, German courts are believed to be impartial and resistant to systematic differences in their interpretation of legal norms. Germany has a civil law system and these legal systems are known to generate decisions in a bureaucratic and objective way. Out of all civil law systems, Germany is its prime example (Langbein, 1985).

But despite of these traditionally held believes, recent scholarship in fact did find system-

---

<sup>6</sup>We leave out all citations to idiosyncratic legal sources, meaning any legal source cited by only one decision. These idiosyncratic legal sources do not offer any information about locations in the context of our model.

atic differences in courts' interpretation of German legal doctrine (Bechtold, Frankenreiter and Klerman, 2019). In areas like press law where German citizens can choose the court they want to file their case with, litigants indeed strategically “shop” for the forums that suit their purpose. In return, courts who wish to establish themselves as a go-to-place in a certain legal area rely on different means to be particularly appealing. Establishing a well known reputation in the tendency of a court's decision making is a successful strategy when selling the own forum – an observation also echoed by recent press reports (Dahlkamp and Schmid, 2014; van Lijnden, 2016).

Jürgens (2014) gathers data on press law caseload at Germany's regional courts (Landesgerichte). He finds that between 2010 and 2012, Berlin (28.67%), Hamburg (22.32%), and Cologne (11.94%) adopt the majority of all legal decisions. The rest of the caseload is shared among other courts. Authors disagree on how to interpret these data. Some argue, that there are systematic tendencies in rule-making favoring litigants who defend themselves against the press at these three places (Dahlkamp and Schmid, 2014; Höch, 2018; Jürgens, 2014, 2016; Kompa, 2012; Sajuntz, 2014; van Lijnden, 2016). The high case load is a result of litigants who – aware of these tendencies – select their courts strategically. In addition to the data from case loads, authors substantiate their claims with a number of case studies. Interestingly, all anecdotal evidence is drawn from either cases with the court in Cologne (Dahlkamp and Schmid, 2014; Jürgens, 2014, 2016) or Hamburg (Höch, 2018; Jürgens, 2014, 2016; Kompa, 2012). Others contend that this interpretation is wrong. The high number of cases in Berlin, Hamburg or Cologne is not a result of forum shopping, but instead the result of a concentration of media companies – including an eco-system of specialized press lawyers who can serve them (Dölling, 2015; Höcker and Brost, 2015). Interestingly, all evidence that is brought forward to refute these claims stems from cases in Berlin (Dölling, 2015). Answering the call for an empirical analysis that “meets the standards of an empirical science (Dölling, 2015, 130, *own translation*)”, we investigate these claims with our measurement model. We thus expect that the courts in Hamburg and Cologne are more likely to adopt decisions that are friendly to litigants and hostile towards the press.

For our model to work, we need to be able to make a core assumption about the judicial

decisions: their circumstances have to be comparable and the relative differences between any two decisions can be mapped on a single dimension; or put differently, the case-space of the decisions we analyze has to be clearly uni-dimensional. We investigate two narrowly defined spaces. First, we select nine recent decisions on privacy infringements through the publication of photos in print media and a respective claim for compensation. Second, we select six decisions on privacy infringement through criminal act allegation in the media and the corresponding claim for injunction. Figure 1 provides an overview over the resulting decision-source matrices  $Y_{ij}$ . Each row represents a judicial decision, each column a legal source. The darker the color, the more often does a decision cite a particular legal source. For the cases on compensation, the left figure displays a number of long vertical lines that indicate a high degree of overlap: These decisions cite similar legal sources – but some more and others less often. In addition, there are a number of legal sources that get cited by only a few judicial decisions. The connectivity matrix for the second application looks quite similar, with some legal sources being standard sources and others that are picked up by a subset of the judicial decisions, only. In conclusion, the respective case-spaces seem to be narrowly defined. The decisions draw on a similar set of legal sources, allowing to meaningfully locate and compare the decisions in the respective uni-dimensional space.

## 5.2 Estimates Of Decision Locations

We use the respective decision-source matrix to estimate a decision’s location. The resulting position estimates are in line with *ex-ante* expectations based on expert knowledge and media reports. Figure 2 depicts the estimated decision locations. The figure displays all decision locations in our samples from Cologne (*Köln*) or Hamburg in red and locations of any other court decisions in blue. A point represents the respective median of the posterior draws. Uncertainty bars around the estimate depict the central 90% credible interval. The box plots on top summarize the distribution of estimated locations from Cologne (*Köln*) and Hamburg *vis-à-vis* other courts. For the decisions related to compensation, we observe two outliers – both from Cologne. However, even though Hamburg and Berlin are also known as friendly towards compensation, the hand-picked decisions

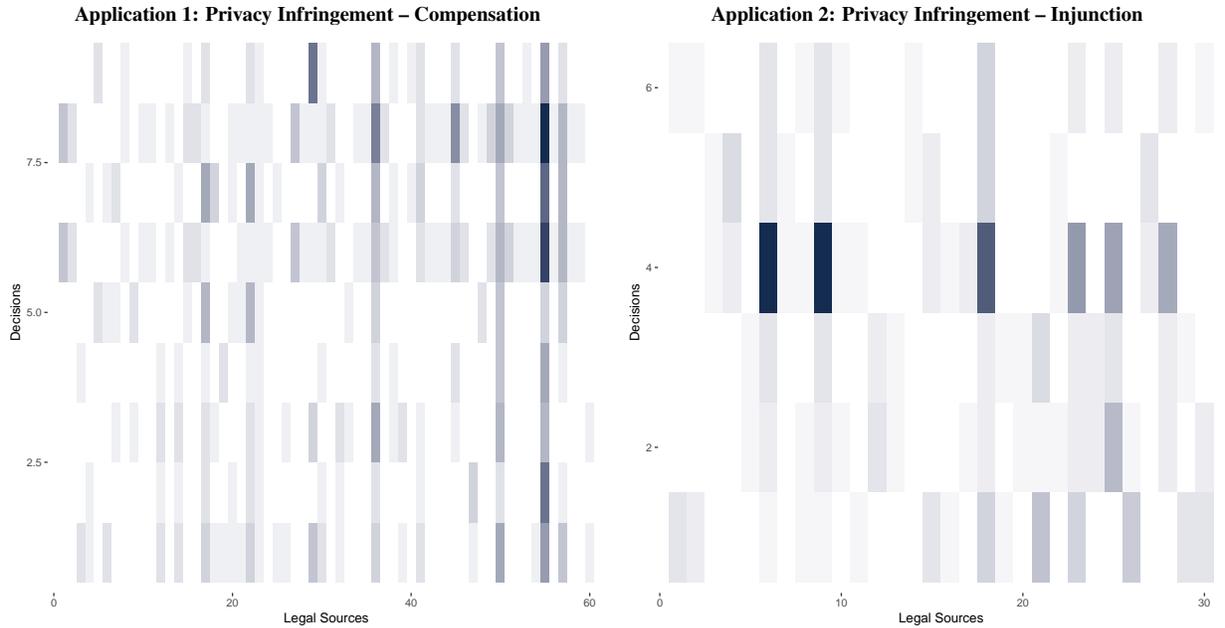


Figure 1: Pre-selected Set of Decisions. Coding of the Decision-Source Matrix on the Basis of Metadata. The Darker the Shading, the More Often a Legal Decision Cites a Legal Source.

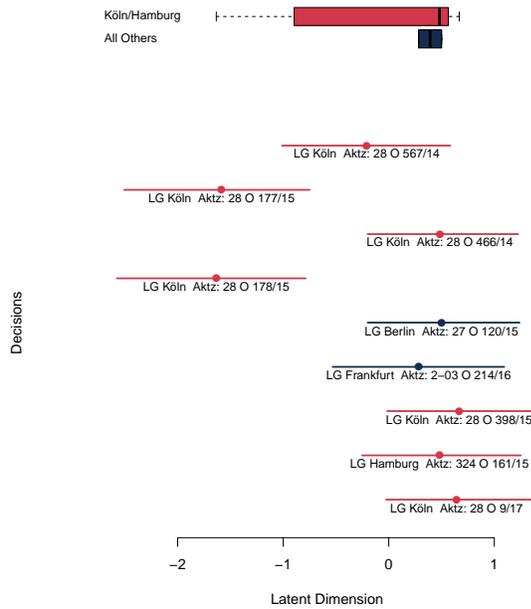
do not show systematic differences. In the decisions related to injunction, the decisions from Cologne and from Hamburg cluster, i.e. they are mapped onto a similar location on the case-space. Heidelberg, the only other court in the sample, is distinctly situated on the right.

In addition to estimates for the location of a written decision, our model also estimates the location of the cited legal sources within the same case-space. Figure 3 shows the estimated locations ( $\hat{\phi}_i$ ) of every cited legal source.

Thus, in contrast to the previous literature we do not only scale decisions but also legal sources in the same case-space. This should facilitate more fine-grained substantive interpretations of the legal argumentation that is developed within those decisions because we type and the frequency of the cited legal sources provide additional information that has not been leveraged before. For our application of compensation cases we see that we could identify a greater variety of legal sources that cluster around  $-1$  and  $.5$  respectively than in our application of the injunction cases which are more spread out and, hence seemingly more informative about the location of the citing decisions.

To sum-up, we find empirical evidence that corresponds to the experts' and journalists' ex-

### Application 1: Privacy Infringement – Compensation



### Application 2: Privacy Infringement – Injunction

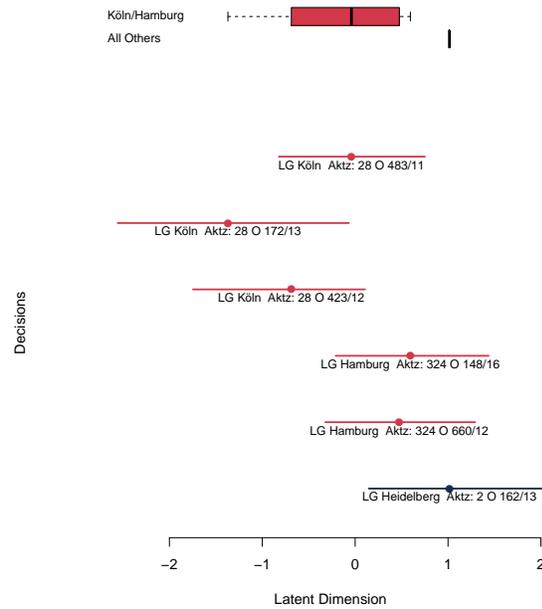


Figure 2: Estimated Locations of Written Decisions ( $\hat{\theta}_i$ ). Pre-selected Set of Decisions. Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

pectations. In the context of press law, lower courts throughout Germany show tendencies in their judgments when it comes to compensation or injunction in privacy infringement cases.

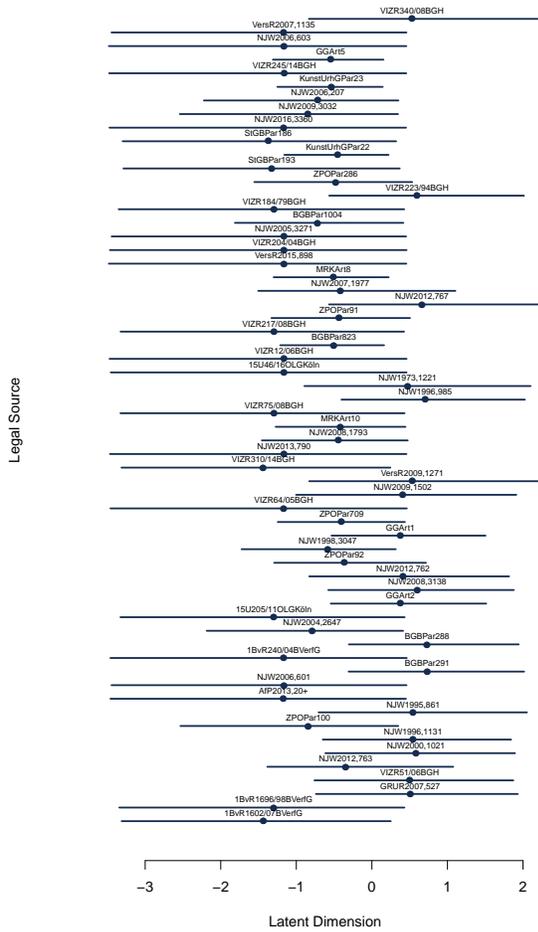
### 5.3 "More-Like-This" Queries Using Titles of Preselected Decisions

So far we investigated the location of decisions using curated collections of judicial decisions to *a priori* define a common space. Can we generalize our approach and automate the case selection above and beyond our pre-selected cases? After all, our ultimate goal is to devise a tool that helps legal practitioners in getting a quick overview over a series of decisions.

We further enlarge the sample of decisions we study with a search engine. Our process is exactly the same as if we used a query field of a search engine in a browser. We “copy” the titles from the two sets of judicial decisions about privacy infringements – the one on compensation and the one on injunction – and then “paste” them into the query field.<sup>7</sup> The resulting top 25 hits are

<sup>7</sup>An example for the collection on compensation cases is decision ‘LG Köln Aktz: 28 O 567/14’ with the following (German) title: “*Unterlassungsanspruch hinsichtlich der Bildnisveröffentlichungen wegen Verletzung des allgemeinen*

Application 1: Privacy Infringement – Compensation



Application 2: Privacy Infringement – Injunction

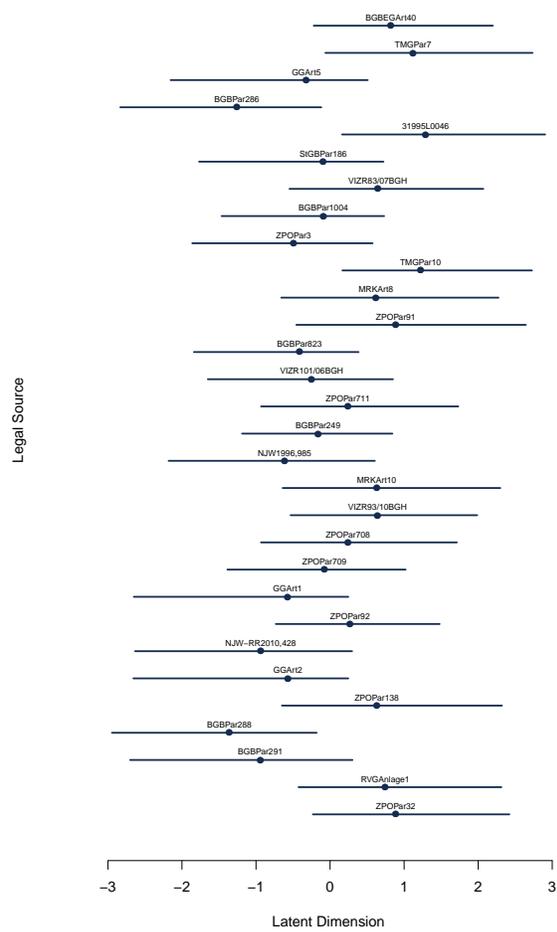


Figure 3: Estimated Locations of Cited Legal Sources ( $\hat{\phi}_j$ ). Pre-selected Set of Decisions. Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

the now enlarged sample.

This is what is happening from a technical point of view.<sup>8</sup> Unlike for the estimation of locations in the case space where we use citation data, we now rely on the actual text corpus of the legal documents instead. The decisions' titles contain the core information in a concise and unequivocal way, making them particularly useful query terms. When looking for similar legal documents in the data base, the algorithm instead considers the complete text to offer a larger base for comparison. In a first step, the data has to be converted into a format that allows computers to calculate the similarity between the query document and all other documents in the data base. All documents in the data base can be represented with one large term-document matrix. This matrix contains the whole vocabulary of the corpus on the first dimension. The second dimension holds the word counts of each document in the corpus.<sup>9</sup> Each document can therefore be represented with a vector  $\vec{v}_i$  that contains all counts for all words. This vector  $\vec{v}_i$  is as long as the size of the vocabulary in the corpus. It turns out that the cosine between two vectors  $\vec{v}_q$  and  $\vec{v}_d$  is a very good measure to calculate the similarity between a query document  $q$  and any other document  $d$  in a corpus. This cosine similarity is calculated as

$$\text{cosine similarity}(q, d) = \frac{\vec{v}_q \cdot \vec{v}_d}{|\vec{v}(q)| \cdot |\vec{v}(d)|}$$

with the numerator being the scalar product between the two vectors  $\vec{v}_q$  and  $\vec{v}_d$ . The denominator standardizes both vectors with the product of their Euclidian lengths.<sup>10</sup> Implementing this measure, we concatenate the titles from the query set to one single query document  $d$  and find

---

*Persönlichkeitsrechts*". For the collection of infringement cases, an example is decision 'LG Heidelberg Aktz: 2 O 162/13' entitled "*Störerhaftung des Betreibers einer Internet-Suchmaschine: Anzeige von Links durch die Suchmaschine zu Internetseiten Dritter mit persönlichkeitsrechtsverletzenden Inhalten*".

<sup>8</sup>For a more in-depth treatment, see [Manning, Raghavan and Schütze \(2009\)](#).

<sup>9</sup>In practice, these counts are being weighted calculating the *term frequency-inverse document frequency (tf-idf)*. This score corrects the pure counts of words in a document with the goal of generating a score that reflects the importance of a word in a document. The score takes into account how often a word occurs relative to all other words and it also corrects for the different lengths of documents.

<sup>10</sup>The Lucene scoring built into the ElasticSearch search engine further refines this measure and allows for more fine grained specifications of search queries. However, these scores are not relevant here. For further information please refer to [https://lucene.apache.org/core/4\\_9\\_0/core/org/apache/lucene/search/similarities/TFIDFSimilarity.html](https://lucene.apache.org/core/4_9_0/core/org/apache/lucene/search/similarities/TFIDFSimilarity.html) (last accessed May 2019)

the most similar judicial decisions available in the data base using the cosine similarity between the tf-idf weighted vectors  $\vec{v}_q$  for the query set and  $\vec{v}_d$  for all others.

On the basis of this similarity measure, we retrieve the 25 decisions that are most similar to our seed decisions, i.e. the set of pre-selected decisions on compensation and infringement. Based on these 25 most (cosine) similar decisions we generate the respective decision-source matrices, again omitting idiosyncratic references (figure 4). For the decisions on compensation, the matrix is quite well-behaved and shows a substantial degree of overlap between those decisions and the cited legal sources. We can clearly identify five legal sources that are being cited by a large number of decisions. In addition, there are legal sources that are being cited only by some decisions, which in the end provide most of the analytic leverage. The connectivity matrix of the second set of written decisions on injunction shows less overlap. Only two legal sources are apparently widely cited.

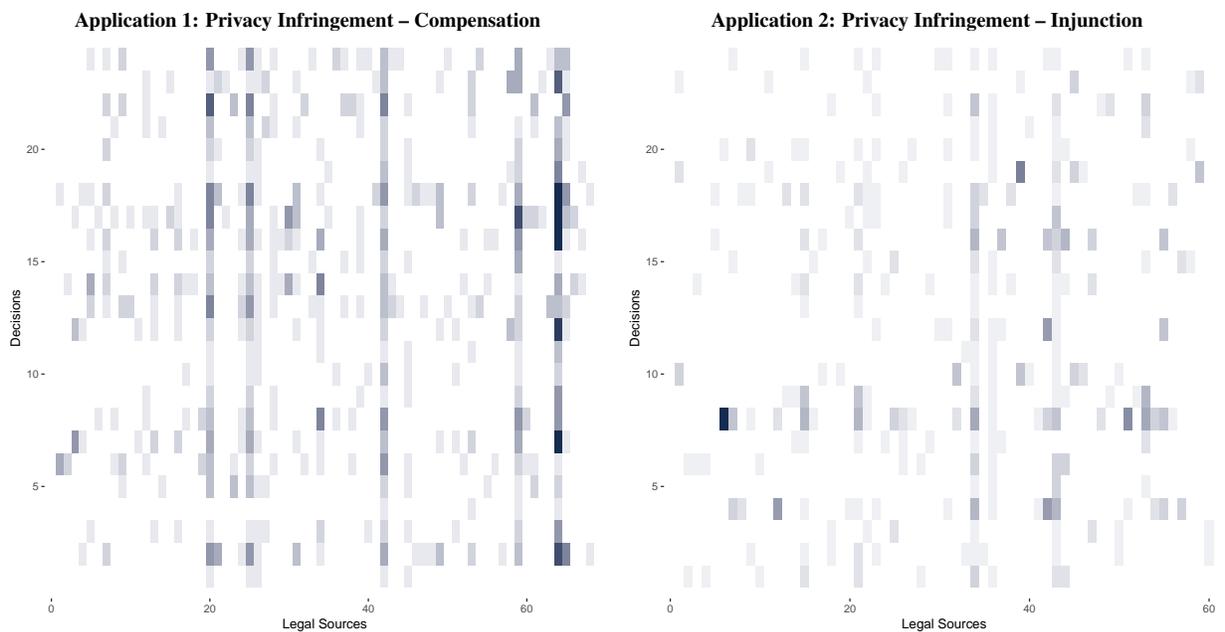


Figure 4: Set of Decisions Using a ‘More-Like-This’ Query. Coding of the Decision-Source Matrix on the Basis of Metadata. The Darker the Shading, the More Often a Legal Decision Cites a Legal Source.

The location estimates in figure 5 reflect the data input. To estimate the decision locations, we use the same setup for the Bayesian estimation as previously. Again, the figure displays the median estimates with a circle, the bars indicate the core 90% credible interval. In the first set of

judicial decisions – the collection on compensation – the courts in Cologne and Hamburg cluster on one end of the spectrum; all other courts can be found on the other end. Decisions on injunctions show more mixed positions. However, as the boxplots on top of the figure show, there is still a trend for the courts of Hamburg and Cologne to be on one end of the spectrum, and all others to be on the other end. Given the size of the credible intervals, though, the locations of many decisions cannot be distinguished on substantive grounds.

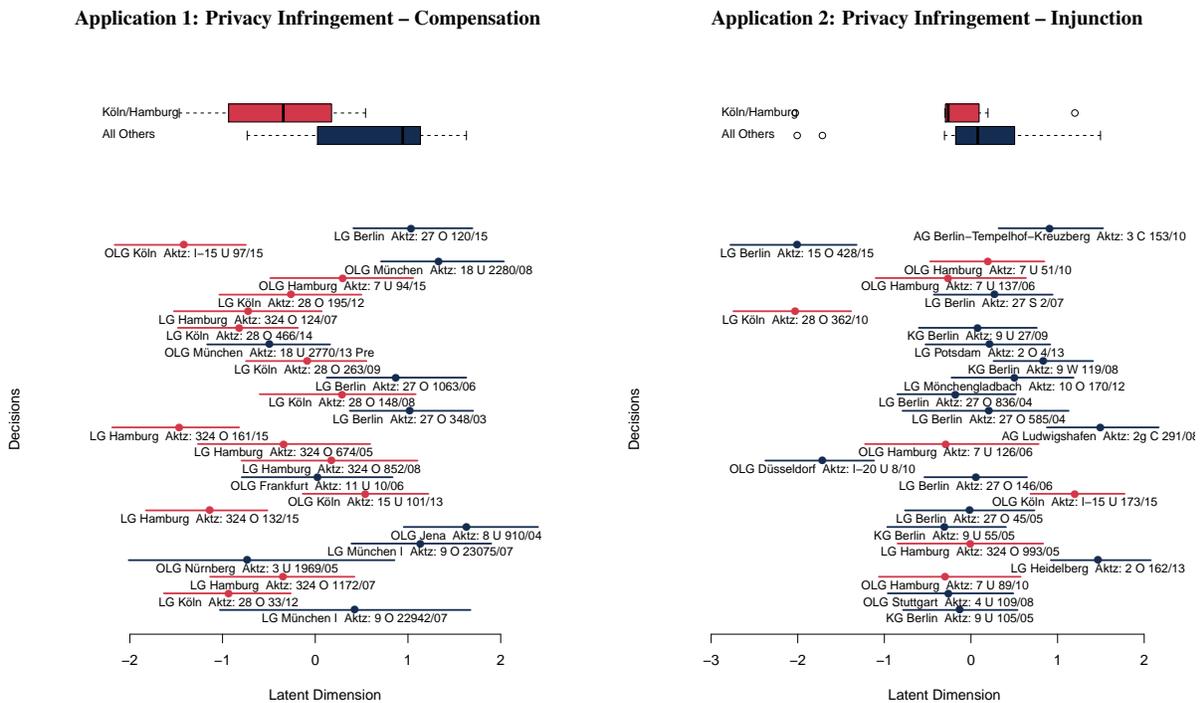


Figure 5: Estimated Locations of Written Decisions ( $\hat{\theta}_i$ ). Set of Decisions Using a ‘More-Like-This’ Query. Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

To sum up, using techniques known from information retrieval help us to automatically enlarge our initial sample of selected seed decisions with similar decisions that can be validity mapped into the same case-space. The estimated locations are in line with anecdotal evidence from media reports and from experts.

## 5.4 Validation and Sensitivity

In the following we assess the validity of our estimates and the sensitivity of our model specification. A first strategy is to leverage a case study to assess the validity of our estimates in several ways. First we use the decision text of two decisions with very similar estimated locations and show that the difference in substance is in fact small. Second, we compare two decisions with different estimated locations and show that the substance is in fact very different as well.

Another strategy to also test the sensitivity of the model specification is to look how the model fits the data. We show below that our model predicts a pattern of citations using the posterior parameter estimates that resembles the true dataset.

### 5.4.1 Case Study to Assess the Validity

To assess the validity of our automated scaling approach we manually compare the substantial content of three decisions of Application 1 (Privacy Infringement / Compensation): two decisions that are scaled at similar positions (LG Hamburg, 324 O 161/15 and LG Köln, 28 O 466/14) and one decision that is scaled at the opposite end (LG Munich, 9 O 23075/07). Figure 6 shows their differences on the latent dimension. To calculate these first differences, we subtract the position estimates of each of the decisions' positions for all posterior draws. The figure displays the respective medians and the core 95% credible interval. Our model predicts a high probability that the position of LG Munich, 9 O 23075/07 is different from LG Hamburg, 324 O 161/15 and LG Köln, 28 O 466/14. In return, the probability that LG Hamburg, 324 O 161/15 and LG Köln, 28 O 466/14 are different from one another is relatively low. If our approach is valid, then similar judicial decisions will derive their legal argument using similar legal sources and the judicial decision that is distinct should rely on different legal sources.

The median estimate of the *Hamburg decision* (LG Hamburg, 324 O 161/15) is to the left of the case-space in Figure 5 (Application 1). The litigant in the Hamburg decision requests a compensation for the repeated publication of pictures of her taken while visiting her hospitalized husband – a famous Formula One driver. The litigant used various legal means to stop the defen-

**Application 1: Privacy Infringement – Compensation**  
 First Differences of Three Decisions

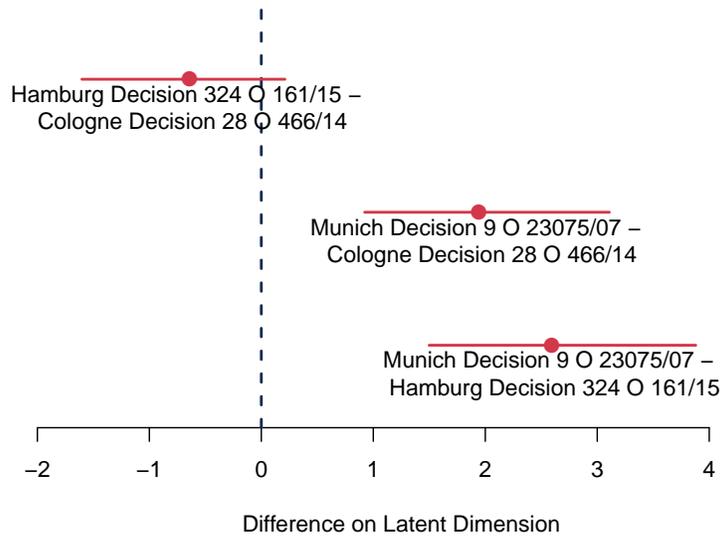


Figure 6: Estimated Differences in the Location ( $\hat{\theta}_i$ ) for Hamburg Decision 324 O 161/15, Cologne Decision 28 O 466/14 and Munich Decision 9 O 23075/07. Points Indicate the Median of the Difference of the Respective Posterior Draws. The Bars Represent the Central 95% Credible Interval of the Difference of the Respective Posterior Draws.

dant from publishing pictures before referring to the *Landgericht* (324 O 161/15, Mn 6).<sup>11</sup> The defendant requests to dismiss the lawsuit arguing, for example, that the contemporary interest in the hospitalization was not limited towards the Formula One driver but would include how the spouse addresses the stroke of fate (324 O 161/15, Mn 27-29).

The district court concluded that publishing the pictures violated the litigant in her general personality rights derived from the German Civil Code (BGB) in connection with the German Constitution (§ 823 I BGB in connection with Art. 2 I and Art. 1 I GG).

According to the court, the publication of a picture does not *per se* violate a person's general personality rights (LG Hamburg, 324 O 161/15, 35). Instead publishing can be justified when documenting contemporary events in line with the German law regulating art and copyright questions (§ 23 I *Kunsturhebergesetz*, [KUG]). Subsequently, the court had to balance the protection of the private sphere of the individual according to the German Constitution (Art. 1 I GG and Art. 2 I GG) and the European Convention on Human Rights (Art. 8 I ECHR) against the freedom of the press to document contemporary events of importance in line with the German Constitution (Art. 5 I GG) and the ECHR (Art. 10 I ECHR).

To justify the litigant's claim, the district court of Hamburg referred to case-law by the German Federal Court (BGH, VI ZR 51/06; VI ZR 272/06). The BGH had argued in the past that the repeated and tenacious publication of pictures can become a very serious infringement of a litigant's privacy and respective violations require financial compensation (BGH, VI ZR 223/94). In particular, the infringement of a person's privacy is wider when media outlets publish pictures compared to textual articles (BGB, VI ZR 230/08). Moreover, the district court of Hamburg highlighted that the litigant had used legal measures against the defendant to hinder the publication of pictures. Hence, the defendant was well aware of the litigant's disagreement to publish pictures showing her in a personal, very exceptional situation. The infringement of the litigant's privacy was not justified and requires the defendant to pay financial compensation. In short, the court of Hamburg heavily relied on case-law published by the BGB.

---

<sup>11</sup>We refer to sections of interest in court decisions using margin numbers (Mn) as shown in the JURIS database.

If our scaling approach is valid, then the citation pattern of the district court of Cologne in the *Cologne decision* (LG Köln, 28 O 466/14) should be similar. This is due to the fact that the median position of the Cologne decision is scaled to left of the case-space in Figure 5 (Application 1), similar to the median position of the Hamburg decision. The estimated location of both decisions are not systematically different from one another in the common case-space.

The litigant in the Cologne decision is an actress who requests an act of omission and a financial compensation for the online and offline publication of an article together with a picture speculating about a second pregnancy of her. The litigant argues that the picture was taken in a private moment and violates her general personality rights. Moreover, it is known that the litigant is anxious to not share information about her personal life. Instead she succeeded with similar legal claims to not publish pictures during her first pregnancy (28 O 466/14, Mn 7). The defendant requests to dismiss the lawsuit arguing in favor of the contemporary importance of the picture which is in the public interest (28 O 466/14, Mn 15-16). These general case characteristics of the Cologne decision are similar to the Hamburg decision. In both decisions the litigants claimed their pictures were taken in private moments violating their personality rights. The defendants rejected claims arguing in favor of the contemporary importance of the pictures.

A closer read of the Cologne decision shows that the district court refers to the same legal norms and similar case-law compared with the Hamburg decision to settle the dispute. Subsequently, the district court of Cologne argues in favor of the litigant's general personality rights derived from the BGB, the GG and the KUG (§§ 1004 and 823 II GCC in connection with Art. 2 I GC and Art. 1 GC as well as §§ 22, 23 KUG). Moreover, the court weights – similar to the district court of Hamburg – the protection of the private sphere (Art. 1 I GG and Art. 2 I GG together with Art. 8 I ECHR) against the freedom of the press (Art. 5 I GG together with Art. 10 I ECHR) referring to the European Convention on Human Rights and the German Constitution (28 O 466/14, Mn 20-23).

In order to finally argue the case, the court in the Cologne decision relies on case-law published by a number of courts and especially the BGH. While not all decisions referred to by the

district court of Cologne were of relevance to the court in Hamburg, both courts heavily relied on case-law published by the BGH. In particular, there is overlap in two BGH rulings repeatedly cited by both courts: VI ZR 223/94 and VI ZR 51/06. These decisions were used to justify the litigant's claim against the defendant.

In sum, it is compelling that the Hamburg decision and the Cologne decision are scaled at similar ends of the case-space in Figure 5 (Application 1). The general case characteristics and the legal outcomes are similar in both cases. Moreover, the legal norms and the case-law used to argue the cases widely overlap.

The median position of (what we call) the *Munich decision* (LG Munich, 9 O 23075/07) is scaled at the opposite end of the case-space in Figure 5 (Application 1) when compared to the median positions of the Hamburg and Cologne decision. If our scaling approach is accurate, we should find that the legal arguments developed in the Munich decision are based on different legal sources than the arguments in the Hamburg and Cologne decision. Moreover, while the cases should address similar scenarios, we might find variation in some case characteristics.

The litigant in the Munich decision -- a famous actress -- requests financial compensation for the publication of secretly taken pictures showing her going for a walk with her newborn. This scenario is comparable to the scenarios in the Hamburg and Cologne decision. However, different to the latter two decisions the defendant in the Munich decision had already given a declaration to refrain from further publications and had payed previous legal fees of the litigant. Nevertheless, the litigant still requested compensation for the pictures already published. The actress argued that the pictures violated her general personality rights, especially as she was in a private moment with her newborn (9 O 23075/07, Mn 3). The defendant requested that the district court rejects the claim. The defendant had already signed a declaration to cease and desist and argued that the litigant is a public figure which is why the pictures were of contemporary interest. The newborn was not recognizable in the pictures (9 O 23075/07, Mn 10-11). The litigant's and the defendant's requests are comparable to the scenarios described in the Hamburg and Cologne decision. Nevertheless, the district court dismissed the litigant's request in the Munich decision and saw no right to receive

compensation (9 O 23075/07, Mn 13-14).

Interestingly, the court argued that the publication of the pictures violated the litigant's rights derived from the German law regulating art and copyright questions (§§ 22 and 23 KUG; 9 O 23075/07, Mn 15). The courts in Hamburg and Cologne presented similar thoughts. Thus, all three decisions seem to be comparable in an appropriate case-space. However, the judges in the Munich decision do not derive a financial compensation from the violation. Instead, the court's line of reasoning is different to the ones presented by the district courts in Hamburg and Cologne. The court in Munich referred to different legal norms than the other two courts.

The financial compensation for a violation of someone's personality rights after publishing a picture is commonly based on Art. 1 and Art. 2 I GG in conjunction with § 823 I BGB. These are the regulations referred to in the Hamburg and Cologne decision. Nevertheless, in the Munich decision the judges refer to decisions by the BGH which the other district courts do not quote. Accordingly, the BGH had outlined that financial compensation requires a "very serious infringement" (9 O 23075/07, Mn 17; own translation) of someone's personality rights. However, the unjustified publication of a picture lowers the legal barriers to receive financial compensation (VI ZR 56/94; VI ZR 255/03). Nevertheless, the district court of Munich concludes that the publication of the actresses' picture does not qualify as serious infringement of her privacy. Instead, the scenario is different to the ones in decisions by the BGH. The legal norm of relevance to the court in Munich is the newly edited § 253 II BGB. This norm allows to grant financial compensation for various physical and psychological violations but — according to the court — is not intended to justify universal compensation. Case-law of courts which had to assess very serious infringements supports this view. Subsequently, in the case at hand the district court of Munich does not see any justification for financial compensation (9 O 23075/07, Mn 16-22).

The district court in Munich faced a scenario similar to the courts in Hamburg and Cologne. Nevertheless, while the case characteristics are comparable the court in the Munich decision argues the case differently compared to the district courts in Hamburg and Cologne. Eventually the court in Cologne derives different legal consequences rejecting the litigant's claim. This is

why it is compelling that the median position of the Munich decision is located at the opposite end to the median positions of the Hamburg and Cologne decision in Figure 5 (Application 1). Subsequently, the comparisons of the three cases supports the validity of our approach to estimate similar locations of the Hamburg and the Cologne decisions, which are very different from the Munich decision.

Figure 7 summarizes similarities and differences between the three decisions assessed throughout the case study. The figure is read from top to bottom. Following the solid line (Hamburg decision), dashed line (Cologne decision) and dotted line (Munich decision) leads to the legal norms and the case-law which are used to develop the judicial arguments over the course of the different decisions. The figure illustrates that our scaling approach allows to differentiate between different nuances in the main body of the decisions, where courts argue the cases. In the initial parts of the decisions there is overlap between the legal norms on which a case is based on. This is not surprising as these norms are presented by the litigant to initiate the cases based on similar scenarios. However, the Hamburg and Cologne decision argue based on similar legal norms and case-law while the Munich decision has nothing in common with the other two decisions. This speaks to the validity of our scaling approach according to which the median scores estimated for the Hamburg and Cologne decision are located at similar ends of our case-space while the Munich decision is placed at the opposite end (Figure 5, Application 1). Moreover, the referrals to different norms and case-law leads to different outcomes favoring either the litigant or the defendant. Finally, the procedural norms to enforce the outcome are listed in Figure 7 for completion. There is some overlap across the decisions referring to the procedural norms. The procedural norms from the *Zivilprozessordnung* (ZPO) do not address substantial, content-related variation in decisions but are commonly referred to at the end of a legal dispute – which is why we did not include them in our qualitative analysis.

All in all, the case study presented here outlines that our automated scaling approach allows to systematically locate similarities and differences in legal arguments across court decisions in a case-space.

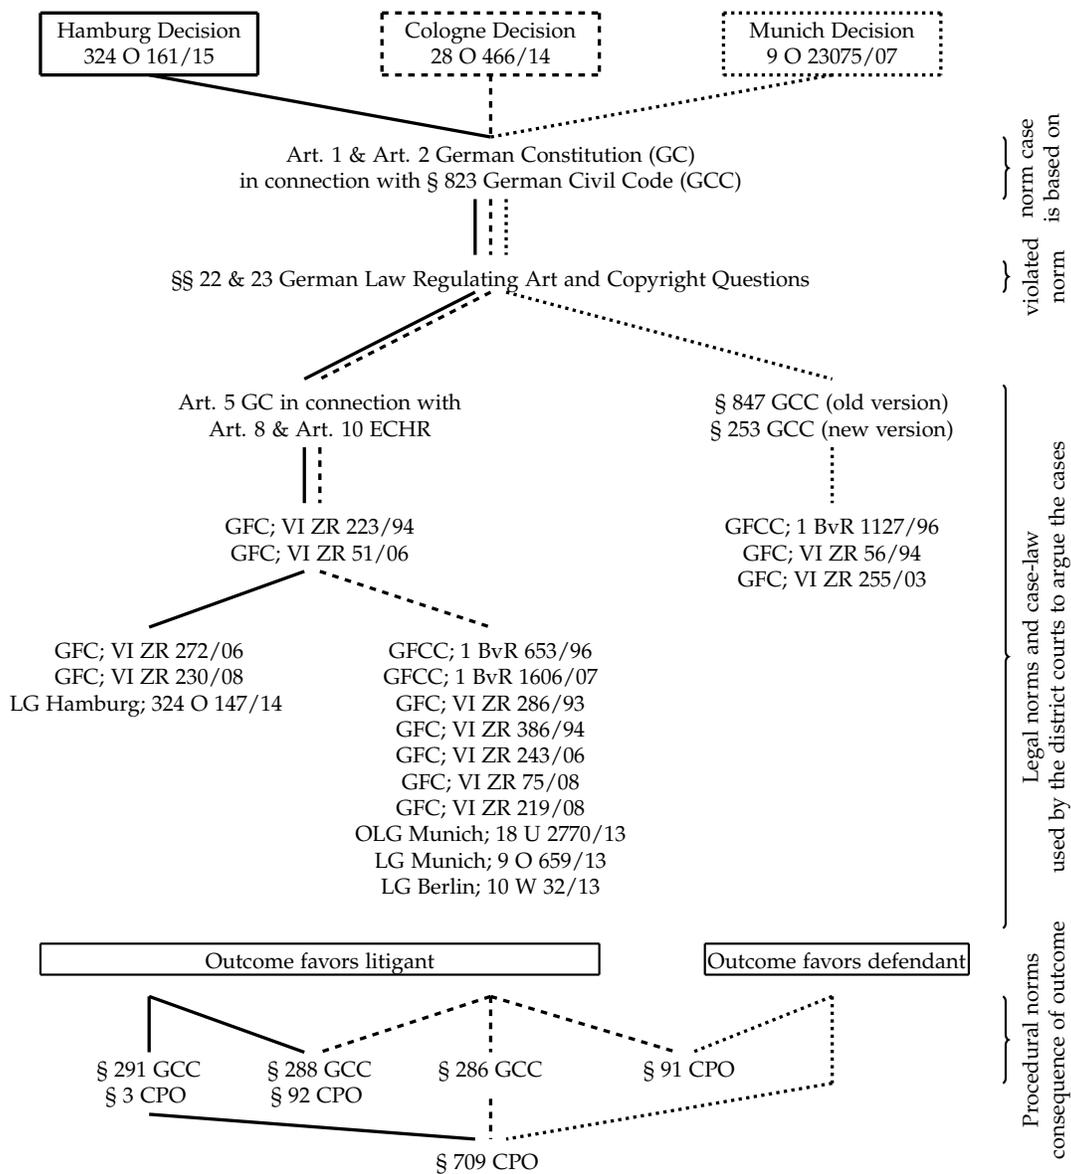


Figure 7: Summary of Legal Argumentation in Three Decisions from Application 1 (Privacy Infringement / Compensation) in Figure 5

### 5.4.2 Assessment of Model Fit using Posterior Predictive Checks

To evaluate the fit of our model, we predict the outcome variable on the basis of the last 50 draws from the posterior sampling chain and plot it against the original data (posterior predictive checks). Figure 8 displays the results for all models. Each figure represents the predicted citation counts  $\hat{y}_{ij}$  for all decision-source pairs as a histogram. The figures display counts horizontally and show their respective squared occurrence vertically.<sup>12</sup> All predictions are in red with credible intervals at the core 80%, 60%, 40% and 20% of the distribution. The red line depicts the median prediction across all models. The dark blue line plots the observed data.

We find that the predictions mostly match the observed data and the models do a good job in predicting the observed counts. For the first application (privacy infringement – compensation), both models do very well as long as the respective decision-source pair is cited frequently. Both model predictions are slightly off for decision-source pairs if observed counts of zero or one or for some artifacts of the observed data — the higher than observed number of citations of decision-source pairs that have four citations. Nevertheless, our model fits the data quite well. Both, our systematic and our stochastic component correctly model the true underlying data generating count process.

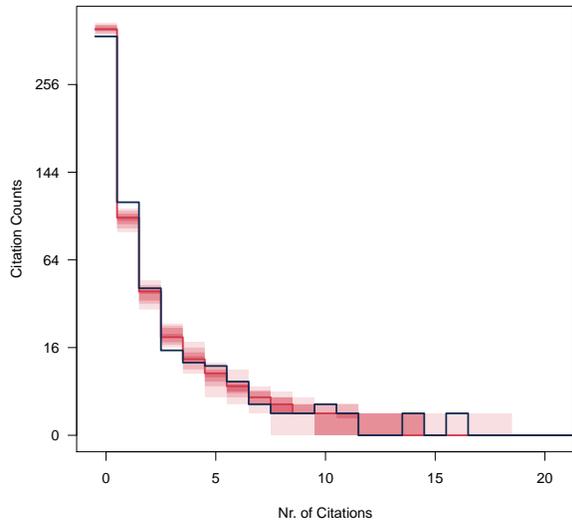
### 5.4.3 Further Validity Test

In addition to the case study and the posterior predictive checks we provide yet another validity test. We leverage the fact that the same decision of a (higher) court can be cited differently. In addition to referring to such a decision directly, some of those decisions are also published in the *Neue Juristische Wochenschrift* (NJW), a German legal magazine. Thus, actually the very same source can be cited in two different ways. From analyzing the meta data from JURIS we do not know that. Thus, our model also does not know that and estimates potentially different locations for the same legal source. A hard validation test is, given that we have two different estimates for the same legal source, to what degree the same legal source is estimated a different location in

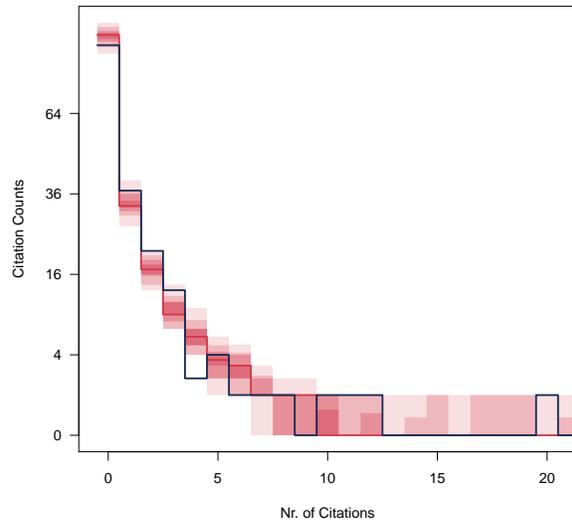
---

<sup>12</sup>In line with [Clark and Lauderdale \(2012\)](#), we display the counts' squared values for better visibility.

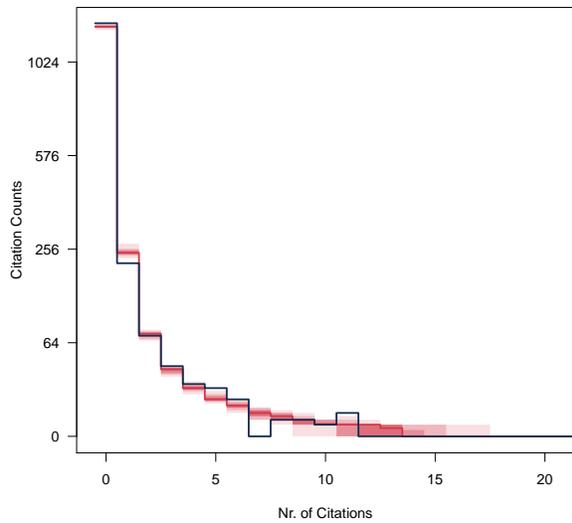
**Application 1: Privacy Infringement – Compensation  
(Curated Sample)**



**Application 2: Privacy Infringement – Injunction  
(Curated Sample)**



**Application 1: Privacy Infringement – Compensation  
(Expanded Sample)**



**Application 2: Privacy Infringement – Injunction  
(Expanded Sample)**

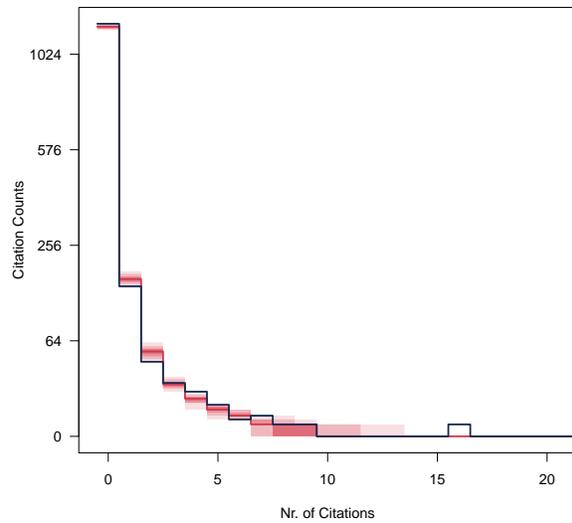


Figure 8: Posterior Predictive Checks. The Figures Visualize the Citation Counts of All Possible Decision Source Pairs in Form of a Histogram. Counts Are Displayed Horizontally and Their Respective Frequencies Is Shown Vertically. Predictions are in Red and stem from the Last 50 Draws of the Sampling Chain with Credible Intervals at 80%, 60%, 40% and 20%. Median Predicted Counts Are Represented with a Red Line, the Observed Data with a Blue Line.

the underlying case-space if the label referring to this source is different. Our first analysis of this shows, however, that we do not find systematically different locations of the same legal source that is referred to in different ways.

## 6 Conclusion

Legal scholars and political scientists are not only interested to understand the decisions judges make, but also seek to analyze the policy implications of written decisions. Previous strategies to compare the locations of decisions in a common case-space rely on voting data of the opinion authors (Clark and Lauderdale, 2010). How can we estimate and compare positions of various actors in a common case-space of court decisions where the individual votes are not published? Answers to this question are of great academic importance in order to explain the variance across cases, judges and courts in common law or civil law systems.

So far no strategy exists that can be applied to scale decisions in a common space when no such voting record exists. The contribution of this paper is that we lay-out a strategy how to estimate and consequently compare positions of various actors in a common case-space of court decisions. Our statistical model relies on leveraging citations networks using the idea that ‘similar’ decisions, i.e. decisions with similar locations in a common space, should cite similar legal sources using similar citation patterns. Our scaling model also locates cited legal sources in the same case-space.

We show that if we pre-select comparable decisions that can be mapped into a common space as long as they have similar citation patterns. We provide so far two applications for appropriate uni-dimensional case-spaces, first privacy infringement cases in combination with either compensation claims or in combination with injunction. Moreover, we show that we can boost an initial seed sample of comparable decisions by ‘more-like-this’ queries using the titles of our pre-selected seed decisions as query terms. The results of both applications are consistent with qualitative accounts from experts and journalists on the phenomenon of *forum shopping*.

However, while we do manage to find evidence in the context of German press law, the main restriction when using this model is the selection of a comparable set of written decisions. Only if we define the set of decisions along very narrow lines – a space that could be ordered along a line by a human reader – we can expect to be successful in estimating valid locations of decisions. This type of validation is currently done and can be hopefully reported soon.

Another limitation of our model is that we ignore the time dimension. Our model is innocent about the development of certain arguments that are characterized by certain citation patterns over time. In fact by pre-selecting our seed decisions we payed attention to that. If there is some time between any two decisions then older ones have obviously a higher likelihood to get cited. This time dimension is not included in the our data-generating process that describes why and how often judges cite particular legal sources.

Locating written lower court decisions in a common space has also important practical implications for law firms. Our model locates every decisions and, hence, also potentially every court somewhere on a latent case-space and allows to identify actual decisions that are anywhere on that continuum. Suppose a lawyer who is in favor of an extreme position in an ongoing case is trying to build an argument. Using our scaling model, she can obtain a quick overview over the tendency in each ruling she selects. This does not only help her identify the locations of those who are in favor of her own opinion. She could also study the arguments of the opposite side to be able to anticipate and preempt them. In addition, the location of actual decisions in a case-space are also very helpful for judges when drafting an decision for a similar case. The comparability of locations of decisions on a latent dimension allows them to quickly locate similar decisions that had similar positions or counter-argue decision that are on the other side of that dimension.

## References

- Barberá, Pablo. 2015. “Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data.” *Political Analysis* 23(1):76–91.
- Bechtold, Stefan, Jens Frankenreiter and Daniel Klerman. 2019. “Forum selling abroad.” *Southern California Law Review* .
- Brouard, Sylvain and Christoph Hönnige. 2017. “Constitutional courts as veto players: Lessons from the United States, France and Germany.” *European Journal of Political Research* 56(3):529–552.
- Carrubba, Cliff, Barry Friedman, Andrew D. Martin and Georg Vanberg. 2012. “Who Controls the Content of Supreme Court Opinions?” *American Journal of Political Science* 56(2):400–412.
- Clark, Tom S. and Benjamin E. Lauderdale. 2010. “Locating Supreme Court Opinions in Doctrine Space.” *American Journal of Political Science* 54(4):871–890.
- Clark, Tom S. and Benjamin E. Lauderdale. 2012. “The Genealogy of Law.” *Political Analysis* 20(3):329–350.
- Clinton, Joshua D., Simon Jackman and Douglas Rivers. 2004. “The Statistical Analysis of Roll Call Data.” *The American Political Science Review* 98(2):355–370.
- Coupette, Corinna and Andreas M. Fleckner. 2018. “Quantitative Rechtswissenschaft Sammlung, Analyse und Kommunikation juristischer Daten.” *JuristenZeitung* 73(8):379–389.
- Dahlkamp, Jürgen and Barbara Schmid. 2014. “„Köln nimmt das alles.“” *Der Spiegel* (42):150–153.
- Dölling, Birger. 2015. “Der fliegende Gerichtsstand im Presserecht – Spielball der Interessen ?” *Neue Juristische Wochenschrift* 43:124–130.
- Dyevre, Arthur. 2019. “The Promise and Pitfalls of Automated Text-Scaling Techniques for the Analysis of Judicial Opinions.” *Working Paper* .
- Epstein, Lee, Andrew D. Martin, Jeffrey A. Segal and Chad Westerland. 2007. “The Judicial Common Space.” *Journal of Law, Economics, & Organization* 23(2):303–325.
- Evans, Michael, Wayne McIntosh, Jimmy Lin and Cynthia Cates. 2007. “Recounting the courts? Applying automated content analysis to enhance empirical legal research.” *Journal of Empirical Legal Studies* 4(4):1007–1039.
- Fowler, James H., Timothy R. Johnson, James F. Spriggs, Sangick Jeon and Paul J. Wahlbeck. 2007. “Network analysis and the law: Measuring the legal importance of precedents at the U.S. Supreme Court.” *Political Analysis* 15(3):324–346.
- Frankenreiter, Jens. 2017. “The Politics of Citations at the ECJ-Policy Preferences of E.U. Member State Governments and the Citation Behavior of Judges at the European Court of Justice.” *Journal of Empirical Legal Studies* 14(4):813–857.

- Gelman, Andrew and Jennifer Hill. 2007. *Data Analysis Using Regression and Multi-level/Hierarchical Models*. Cambridge: Cambridge University Press Cambridge.
- Hall, Mark A and Ronald F Wright. 2008. "Systematic Content Analysis of Judicial Opinions." *California Law Review* 96(1):63–122.
- Höch, Dominik. 2018. "Alles neu, alles anders? Das BVerfG nimmt sich der Gerichtspraxis im Presserecht an."  
**URL:** [www.zpoblog.de/?p=6790](http://www.zpoblog.de/?p=6790)
- Höcker, Ralf and Lucas Brost. 2015. "Kompetenz zahlt sich aus - Zur Notwendigkeit des fliegenden Gerichtsstandes im Nischenrechtsgebiet Presserecht." *IP - Rechtsberater* (6):138–141.
- Hönnige, Christoph. 2007. *Verfassungsgericht, Regierung und Opposition. Die vergleichende Analyse eines Spannungsdreiecks*. Wiesbaden: VS Verlag für Sozialwissenschaft.
- Hönnige, Christoph. 2009. "The Electoral Connection: How the Pivotal Judge Affects Oppositional Success at European Constitutional Courts." *West European Politics* 32(5):963–984.
- Jackman, Simon. 2001. "Multidimensional Analysis of Roll Call Data via Bayesian Simulation: Identification, Estimation, Inference, and Model Checking." *Political Analysis* 9(3):227–241.
- Jackman, Simon. 2004. "Bayesian Analysis for Political Research." *Annual Review of Political Science* 7:483–505.
- Jackman, Simon. 2009. *Bayesian Analysis for the Social Sciences*. New York: Wiley.
- Jacob, Marc A. 2014. *Precedents and Case-Based Reasoning in the European Court of Justice: Unfinished Business*. Cambridge: Cambridge University Press.
- Jürgens, Uwe. 2014. "Abgestürzte Gerichtsstände – Der fliegende Gerichtsstand im Presserecht." *Neue Juristische Wochenschrift* 42:3061–67.
- Jürgens, Uwe. 2016. "Turbulenzen im Presseprozessrecht. Der rechtswidrige Standardverzicht auf mündliche Verhandlungen im einstweiligen Rechtsschutz." *Kommunikation und Recht* (1):7–12.
- Kelemen, Katalin. 2013. "Dissenting Opinions in Constitutional Courts." *German Law Journal* 14(8):1345–1371.
- Klerman, Daniel and Greg Reilly. 2016. "Forum Selling." *Southern California Law Review* 89(2):241–316.
- Kompa, Markus. 2012. "Experiment zum fliegenden Gerichtsstand."
- Kornhauser, Lewis A. 1992a. "Modeling Collegial Courts I: Path-Dependency." *International Review of Law and Economics* 12(2):169–185.
- Kornhauser, Lewis A. 1992b. "Modeling Collegial Courts II: Legal Doctrine." *Journal of Law, Economics, & Organization* 8(3):441–470.

- Landa, Dimitri and Jeffrey R Lax. 2008. "Disagreements on Collegial Courts: A Case-Space Approach." *Journal of Constitutional Law* 102:305–329.
- Langbein, John H. 1985. "The German Advantage in Civil Procedure." *The University of Chicago Law Review* 52(4):823.  
**URL:** [http://digitalcommons.law.yale.edu/fss\\_papers/536](http://digitalcommons.law.yale.edu/fss_papers/536)
- Lax, Jeffrey R. 2011. "The New Judicial Politics of Legal Doctrine." *Annual Review of Political Science* 14(1):131–157.
- Lax, Jeffrey R. 2012. "Political constraints on legal doctrine: How hierarchy shapes the law." *Journal of Politics* 74(3):765–781.
- Manning, Christopher D., Prabhakar Raghavan and Hinrich Schütze. 2009. Scoring, term weighting, and the vector space model. In *Introduction to Information Retrieval*. Cambridge: Cambridge University Press chapter 6, pp. 118–132.  
**URL:** <http://nlp.stanford.edu/IR-book/pdf/06vect.pdf>
- Martin, Andrew D. and Kevin M. Quinn. 2002. "Dynamic Ideal Point Estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953-1999." *Political Analysis* 10(2):134–153.
- McGuire, Kevin T. and Georg Vanberg. 2005. "Mapping the Policies of the U.S. Supreme Court: Data, Opinions, and Constitutional Law." *Working Paper* .
- Poole, Keith T. and Howard Rosenthal. 2007. *Ideology and Congress*. Transaction Publishers.
- Raffaelli, Rosa. 2012. "Dissenting opinions in the Supreme Courts of the Member States." *European Parliament: Directorate General for Internal Policies* (PE 462.470).
- Sajuntz, Sascha. 2014. "Die Entwicklung des Presse- und Äußerungsrechts in den Jahren 2012/2013." *Neue Juristische Wochenschrift* pp. 25–30.
- Sternberg, Sebastian, Thomas Gschwend, Caroline E. Wittig and Benjamin G. Engst. 2015. "Zum Einfluss der öffentlichen Meinung auf Entscheidungen des Bundesverfassungsgerichts: Eine Analyse von abstrakten Normenkontrollen sowie Bund-Länder-Streitigkeiten 1974 - 2010." *Politische Vierteljahresschrift* 56(4):570–598.
- van Lijnden, Constantin. 2016. "Fliegende Richter." *Die Zeit* (28):1–5.
- Whalen, Ryan. 2016. "Legal Networks: The Promises and Challenges of Legal Network Analysis." *Michigan State Law Review* (539):539–566.
- Winkels, Radboud, Jelle de Ruyter and Henryk Kroese. 2011. "Determining Authority of Dutch Case Law." *Legal Knowledge and Information Systems* 235:103–112.