

Scaling Court Decisions with Citation Networks*

CHRISTIAN ARNOLD[†]
Cardiff University

BENJAMIN G. ENGST[‡]
University of Mannheim

THOMAS GSCHWEND[§]
University of Mannheim

ABSTRACT

To compare court decisions in a systematic way, it is typically necessary to first read these decisions and then apply legal methods to them. While measurement models exist that support analysts in this manual labour, they usually build on judges' individual voting records and can therefore only be helpful in legal systems where this information exists. This paper proposes a new method that offers the same level of support, but for a more general set of cases. Relying on the readily available references in the courts' decisions, we propose a latent-variable model that measures the latent position of court decisions in a common case-space on the basis of citation networks. We showcase our model in the context of forum-shopping and forum-selling of Germany's lower courts.

*NOTE: Previous versions of this paper have been presented at the *Second Conference on Empirical Legal Studies in Europe*, Leuven, 2018 and the *Conference on Data Science and Law*, Zurich, 2019. *JURIS* generously supported this research by allowing access to their data. We thank Pablo Barberá, Jens Frankenreiter, J. Andrew Harris, Kevin Quinn, Nils Schaks, Philip Schroeder and Alexander Tischbirek for helpful comments. We also thank Felicia Riethmüller for her insightful and instrumental research assistance. Arnold is grateful for support through Cardiff's "Darlithwyr Disglair Development Scheme". Engst and Gschwend acknowledge financial support for SFB 884 (project C4) at the University of Mannheim from the *German Research Foundation* (DFG). All data and computer code necessary to replicate the results in this analysis will be made publicly available on publication.

[†]Senior Lecturer, School of Law and Politics, Cardiff University, Museum Avenue Cardiff, CF10 3AX. (arnoldc6@cardiff.ac.uk).

[‡]PostDoc, Collaborative Research Center 884, University of Mannheim, B6, 30-32, D-68131, Mannheim, Germany (engst@uni-mannheim.de).

[§]Professor, Department of Political Science, University of Mannheim, A5, 6, D-68131, Mannheim, Germany (gschwend@uni-mannheim.de).

1 The Promise of Scaling Lower Court Decisions

How similar is the legal reasoning in comparable court decisions? The answer to this question is of key concern to everyone with an interest in law—scholars or practitioner alike. Some might want to quickly identify particularly controversial decisions. Others seek to get a more comprehensive overview over a large set of decisions, for example to understand patterns across different courts or the development of a legal doctrine over time. This requires legal methods and a close reading of each decision which is realistic for a limited set of decisions. Existing approaches for comparing decisions at large scale mostly rely on data on the voting behavior of the respective judges in a court (e.g. Clark and Lauderdale, 2010). It turns out, however, that in cross-national comparison judges do not always take a vote—and if they do, courts only rarely publish the individual votes (Kelemen, 2013; Raffaelli, 2012). Existing methods only allow to compare courts’ decisions in some countries (e.g. Martin and Quinn, 2002; Hanretty, 2012*a,b*).

We introduce a scaling model that estimates the location of court decisions in a common case-space. Instead of published votes, the model relies on citations.¹ How often does a decision refer to a particular legal source? Similar to Clark and Lauderdale (2010), we assume that the closer the decision to a legal source in a common case-space—and hence the more amenable a source to the legal reasoning in the decision—the more likely is a ‘panel of judges’ (or simply a ‘court’) to refer to this legal source. The main advantage of our approach is that citations are commonly available in every legal system at all levels because judges need these references to justify their reasoning.

For our model, we first pre-select decisions that actually can be compared on legal grounds. While curating a concise set of legal documents can of course be done manually, we show how to help scale the human effort with algorithms from information retrieval. We assess data from the legal data base *Juris*,² and investigate systematic tendencies in judicial decision making in an unlikely case. Ger-

¹We employ a technical definition of the term *citation* and use it to describe any reference to other legal sources, be it other laws, or the reference to other cases or legal literature.

²Court 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.

many has a civil law system that is renown for a particularly impartial and objective way of creating legal decisions (Langbein, 1985). However, a recent debate on ‘forum-selling’ and ‘forum-shopping’ at Germany’s *Landgerichte* (district courts) suggests that there are areas where courts are systematically biased (Klerman and Reilly, 2016; Bechtold, Frankenreiter and Klerman, 2019). Studying cases from press law and antitrust, we indeed corroborate these findings at Germany’s lower courts in Cologne (Köln) and Hamburg for press law, as well as Cologne, Dortmund and Mannheim for antitrust.

We contribute to the literature in three ways. First, we introduce a new model that allows scholars and legal practitioners to systematically compare decisions that share a common case-space on the basis of their references—which is data that is widely available. Second, we showcase how to scale measurement efforts when analyzing decisions from a large legal corpus of lower court decisions. Tools from information retrieval allow to identify decisions that lie in the same common case-space and that can thus be compared to one another. Third, the substantive application of our new scaling model informs recent debates on ‘forum-selling’ and ‘forum-shopping’ because we show that some courts are systematically more plaintiff friendly than other courts. Our model helps investigate systematic biases in court decisions comparatively and at large scale.

The paper proceeds as follows. We begin with surveying existing approaches that are capable of mapping court decisions in a common case-space. The subsequent section explains why citations leave a trace about a court’s tendency in rule making. We then translate this reasoning into a statistical model, and finally study ‘forum-selling’ in Germany’s lower courts in press law and antitrust. A final section concludes.

2 Measuring the Content of Court Decisions

Measuring the relative legal positions of court decisions implicitly accepts the notion that the decisions reside on a continuum: cases that are clearly within the bounds of the law on one end of a spectrum, cases that clearly violate it on the other end, and all others somewhere in between. This spatial notion has been formalized into an analytical framework originally developed to study the influence

of politics in high courts such as the U.S. Supreme Court (for an overview e.g. Clark and Lauderdale, 2010; Lax, 2011). But while the resulting case-space model (Kornhauser, 1992*a,b*; Landa and Lax, 2008; Lax, 2011, 2012) offers a rigorous analytical framework, one of its most significant limitations is an empirical one. So far there are no tools to easily and reliably map a large number of real-world cases into such a common case-space.

Existing empirical strategies offer room for improvement. Broad ideological categories (e.g., liberal-conservative) are of limited use, since they fail to honor the respective legal context. Nevertheless, some even equate the latent political—not the legal—position of a politically nominated (median) judge with a court’s resulting decisions (e.g. Brouard and Hönnige, 2017; Carrubba et al., 2012; Hönnige, 2009; Sternberg et al., 2015). The judge’s displayed political position that is taken as a proxy for her decisions might not even be her own, but can be inherited from other actors who nominated the (median) judge such as parties, senators or presidents (e.g. Epstein et al., 2007; Hönnige, 2009). An alternative approach is to closely analyze the decisions’ content. Law scholars embraced content analysis methods a while ago (Hall and Wright, 2008). However, when hand-coding potentially a large number of written decisions, reliability can indeed be an issue. In addition, since the task is labour intensive, it is hard to scale—particularly so when time and resources are scarce.

There are various ways in which scholars can make the most of computers when they want to locate decisions in a common case-space at large scale. Of course, the words of decisions reflect their meaning. Existing approaches pay close attention to studying the language of court decisions and have been exploring supervised and unsupervised text-scaling (e.g. Dyevre, 2019; Evans et al., 2007; Jakab, Dyevre and Itzcovich, 2017; McGuire and Vanberg, 2005) or dictionary methods (e.g. Owens and Wedeking, 2011, 2012). In the light of the specific requirements of legal terminology, further research in this direction promises important progress. Another option is to rely on the legal sources courts cite in their decisions (Whalen, 2016). The resulting citation networks allow for example to uncover the relevance of a court decision (Coupette and Fleckner, 2018; Fowler et al., 2007; Petersen and Towfigh, 2017; Winkels, de Ruyter and Kroese, 2011). Citation patterns have also been used to show that judges chose legal sources neither randomly nor independently of their preferences. Frankenreiter (2017) exploits the institutional setting at the ECJ where—unlike in other courts—two

opinions are drafted, one by the *Advocate 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. Finally, we are not the first to use citation patterns for explicitly mapping decisions into a common case-space. Clark and Lauderdale (2010) analyze search and seizure cases and freedom of religion opinions before the US Supreme Court between 1953 and 2006. But while Clark and Lauderdale have to consider the judges' voting behavior to estimate valid positions of opinions, we show how to do without such data. Given that in cross-national comparison the highest courts rarely, or never, publish judges' votes (Kelemen, 2013; Raffaelli, 2012), we believe that our approach offers a welcome contribution to the existing toolkit for locating decisions in a common case-space.

In short, current approaches that locate decisions in a common case-space too often work with crude proxies. There are unexplored potentials in going beyond the human effort of reading and coding each single decision. While first work exists that makes use of the information in citation patterns, there is yet no model that allows to map a court decision in a common case-space in a convenient way.

3 Locating Decisions Using Citation Data

The text of a written decision serves the same purpose in any legal system. It provides an argument why a certain case is decided in the particular way it is and not differently. Judges refer to legal sources to bolster their argument. The sources a court is citing in a decision are not only a technical requirement; citations may also reveal legal preferences. Ultimately, the paper seeks to provide nuanced measurements that allow to scale court decisions—which could for example be used to tell a 'hard' from a 'soft' verdict. To explain how to quantify this statement, we first lay a conceptual foundation and take a closer look at the case-space framework. We will then understand why different legal preferences express themselves in the citations of a court decision and finally consider the consequences arising from this observation.

3.1 The Common Case-Space

The case-space model comes with a number of core concepts (Cameron and Kornhauser, 2017*a,b*). A legal case can be defined as a “concrete, fact-ridden dispute between two (or perhaps more) parties (Cameron and Kornhauser, 2017*a*, 2)”. The court has to resolve the dispute—and it does so by applying the law as a rule to the legal facts of the case based on the available evidence. All possible cases on the same subject matter reside on a one-dimensional common case-space. For example, if the police catches car drivers at a certain velocity, it is possible to map all cases into this common case-space on the basis of their speed. For the court to be able to take a decision, it requires a rule that tells the court when the defendant has to be convicted. In the example, there needs to be a law that defines the maximum velocity. A driver below that threshold will not be convicted, a driver above that speed will be.

Turning facts into a disposition on the basis of a rule seems straightforward in the simplified speeding example. And if adjudicating were a simple matter of applying a rule to clear facts, an algorithm would be qualified enough to sit on the bench. The legal reality, however, is far more complex. There are two processes that describe how a case based on available evidence can be mapped into a case-space to arrive at a disposition. They are not strictly deterministic and require human judgement (Cameron and Kornhauser, 2017*a*). First, judges have to translate the available evidence into legal facts. We expect that courts do not differ systematically in their translation from evidence into legal facts. Second, they need to select the appropriate rule and apply it to the identified legal facts. Typically, a court has to evaluate legal facts with a rule on more than one single issue to arrive at a verdict, which is why judges ultimately need to weigh and aggregate all issues when they finally speak the law.

When courts evaluate facts on the basis of a rule—be it laws, existing precedent, previous opinions on similar cases or legal scholarship—they have to provide reasons to legitimize their judgement (Charlotin, 2017) and signal consistency of their reasoning (Lupu and Voeten, 2012) by citing particular legal sources. The set of citations in a decision is indicative of the court’s opinion (Clark and Lauderdale, 2010; Choi and Gulati, 2008). Two courts who apply the same set of rules to assess the

same legal facts will come to the same conclusion citing similar legal sources. If the judges decide a case differently, they are likely to apply different rules to the legal facts and hence will refer to different legal sources.³ Therefore, the legal sources in a set of decisions of the same case-space hold the key for locating the decisions in that common case-space.

3.2 How Do Courts Cite?

The legal sources a court refers to signal important information about the kind of legal argument and the framing of a decision a court is advancing in a decision (Alschner and Charlotin, 2018; Clark and Lauderdale, 2012; Charlotin, 2017; Choi and Gulati, 2008). How do judges pick these legal sources? We argue that there are three different behavioral mechanisms that drive what and how often a source is referred to in a court decision. First, studying citation practices of federal appellate court judges in the US, Choi and Gulati (2008, p.91) argue that the substance matter dictates an authoritative core set of legal sources any judge would mention, for instance, binding rules and norms or previous decisions on similar cases. This argument travels across different legal traditions and also holds for civil-law systems (e.g., Choi and Gulati, 2008; Clark and Lauderdale, 2010). Those *core legal sources* define the legal substance of what is at stake. All court decisions in a common case-space refer to these core legal sources with a similar frequency. Second, there are legal sources we call *idiosyncratic legal sources* because they appear in only one written decision. A court's decision may refer to idiosyncratic legal sources to account for the characteristics of a particular case. Citations of such idiosyncratic legal sources cannot be informative to determine the relative locations among similar decisions within a common case-space. Finally, there are *informative legal sources* that help us estimate the relative locations of decisions within a common case-space. Decisions refer to these sources with a different frequency which reflects their location in the common case-space. In short, we distinguish between three different types of what Posner (2000, p.384) calls "informational" citations. Among these three different citation mechanisms, only the latter one using informative legal sources allows to identify the relative location of a decision in the common case-space.

³Alternatively, judges decide a case differently if they translate the available evidence into different legal facts. It is rather unlikely that this is done without appropriate references if the cases are similar enough to lie within the same case-space.

Many studies of citation patterns find that a decision is more likely to refer to legal sources that are in line with the court’s reasoning (e.g., Choi and Gulati, 2008; Clark and Lauderdale, 2010). Sources that reflect the spirit of the decision help make a legally sound argument (Charlotin, 2017, p.282), which is why the court is likely to quote these legal sources. A court tends to refrain from citing legal sources that are not supportive of the advanced legal argument, because they generate cognitive dissonance and, eventually, cause extra effort when justifying the court’s decision (Charlotin, 2017; Posner, 2000). The court refers to dissenting legal sources—if at all—only in passing and will mention them strategically (Alschner and Charlotin, 2018; Lupu and Voeten, 2012; Lupu and Fowler, 2013). Courts may distinguish their legal argument from existing ones to show why a certain legal belief is not binding or valid by entertaining a potential counterargument or introducing a more nuanced view. Similar to the analysis of precedent in decisions of the ECJ (Jacob, 2014), these negative citations are used to bolster the credibility of the decision, demonstrating the argument’s coherence or fending-off potential revisions from higher courts. In addition, judges try to minimize their chances of reversal by a higher court and to foster their own reputation (Choi and Gulati, 2008; Jacob, 2014).

In result, given limited time and resources, courts are more likely to refer to a legal source in line with the legal reasoning of the decision than a legal source that runs against it. Thus, courts are not only more likely to refer to supportive legal sources in their decisions, they will also refer to them more frequently.⁴

3.3 Implications and Features

To illustrate the implications of the model, we introduce a simple visualization in Figure 1. We begin with the left panel where two decisions D_1 and D_2 in a common case-space refer to the informative legal sources S_1 and S_2 . The thickness of the lines corresponds to how often the two decisions are referring to the sources. The decision D_1 is referring to the source S_1 quite heavily which is why their locations are close to one another. In contrast, the decision D_2 is referring to the two sources S_1 and S_2 to a similar degree and its location is therefore more central. Note that D_2 is not completely in the

⁴Our work bears similar assumptions about the citation process as for example Clark and Lauderdale (2012) who also analyze citation counts.

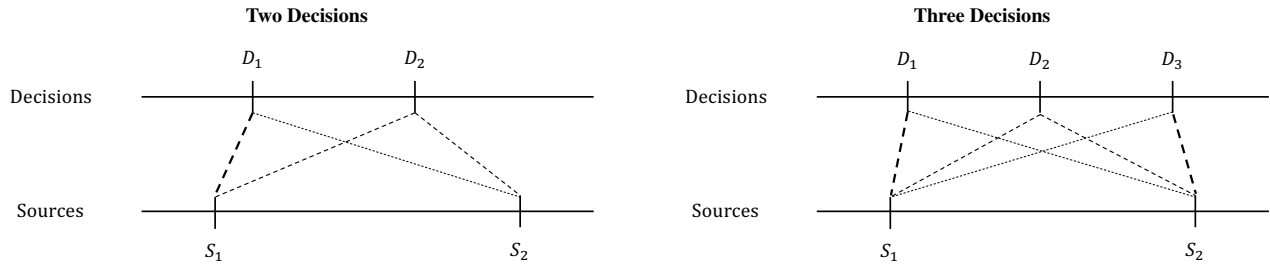


Figure 1: Locations of Decisions D_1 , D_2 and D_3 in the Same Case-Space Given How Often They Refer to Sources S_1 and S_2 . The Thickness of the Lines Corresponds to the Number of Citations. A Third Decision in the Sample Changes the Absolute and Relative Locations of D_1 and D_2 .

middle between the two sources, because of the way in which D_1 is referring to the two sources. The location of the decisions D_1 and D_2 changes when we introduce a third decision D_3 that is referring to S_2 and S_1 to the same degree as D_1 is referring to S_1 and S_2 . The symmetry of citations now enforces the symmetry of the locations. Both of the previous decisions D_1 and even more so D_2 change their location with regards to S_1 and S_2 when adding D_3 .⁵

In addition, adding D_3 also changes the relative locations in the common case-space. Now D_1 and D_2 do not represent the minimum and maximum of the underlying common case-space any longer. Instead, the extreme locations of the new space are defined by D_1 and D_3 . If we define the decisions' common case-space to have a standardized metric—for example z-scores—then the numerical value from D_2 would change to a much more central location once the decision D_3 is included.

In contrast to previous research (e.g. Clark and Lauderdale, 2010) we consider all legal sources and do not exclude any sources, such as procedural legal sources. Also, with our proposed approach it is not necessary to *ex-ante* distinguish between positive and negative citations and between different types of legal sources to estimate the location of lower court decisions in an appropriate common case-space. Positive and negative citations are already accounted for with different expected frequencies in the data generating process, since negative citations occur less often than positive citations. Furthermore, irrespective of the type of the legal sources, the frequency with which lower courts refer to certain legal sources—substantive or procedural—will always be informative because it reflects a

⁵This 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.

choice of the author of a decision.

4 A Case-Space Estimator for Decisions and Cited Legal Sources

Now that we developed an understanding for how courts refer to legal sources when writing decisions, we translate these insights into an appropriate measurement model. The data we observe consists of a $n \times m$ decision-source matrix \mathbf{Y} of citation counts, i.e., y_{ij} represents how often a court's decision $i \in \{1, \dots, n\}$ is referring to legal source $j \in \{1, \dots, m\}$. We use a Poisson distribution as a typical probability model for such citation count data. The systematic component of our model rests on three assumptions. (1) Each decision has a fixed location along a uni-dimensional case-space. (2) A written decision will refer to a legal source more often the closer their locations in this common case-space.⁶ (3) Each decision has a positive probability to refer to any legal source. We express the probability of observing any particular distribution of legal source citations as

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

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

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 .⁷ The parameter γ captures the overall sensitivity of this difference in the respective common case-space. The parameters α_j and β_i explicitly capture the idiosyncrasies of citation counts. The parameter α_j expresses the authority of a legal source j . Some legal sources are by default cited more often than others, simply because they are on average more relevant. Similarly, β_i captures decision-specific differences. Some decisions refer to on average more legal sources than others, for example because they are longer.

Identification is a core concern when estimating latent variable models. These models typically

⁶Coupette and Fleckner (2018, p.384) make the same assumption when laying-out the process how German courts cite legal sources.

⁷This is similar to the strategy of Clark and Lauderdale (2010), but our systematic component employs a richer parametrization (following Barberá, 2015).

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 of these concerns. First, the model has to be anchored in the parameter space so that it does not float around. Second, 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, third, symmetrically rotate the model in the parameter space by inverting all parameters.

Bayesian estimators address identification challenges by specifying prior distributions (Gelman and Hill, 2007; Jackman, 2009). Following Barberá (2015), we solve local identification of our model with a standard normal distribution for θ_i and with a normal distribution with a mean at 0 and standard deviation σ_α for α_j .

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

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

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

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

Global identification is more challenging. Political scientists who estimate locations of political actors and legislative proposals in a common space face a similar problem when they repurpose item response theory (e.g., Clinton, Jackman and Rivers, 2004; Poole and Rosenthal, 2007). They address rotational invariance 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). 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. However, fixing certain decisions or legal sources is not possible in our context. *Ex-ante*, an analyst is neither certain about a clear location of certain decisions nor about 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.

The other option to address rotational invariance is to not identify a model globally at all. Since the model can flip, the respective posterior distribution can have two modes. While in theory the sampler could visit both sides, in practice this concern often turns out not to be an issue as long as there is a reasonably large number of latent positions whose locations can be reasonably well distinguished (Jackman, 2001, 2004). 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). Barberá (2015) follows this latter approach in his original implementation of the model we are building on here. We find, however, that in our application a sampler can indeed flip sides which is why we have to resort to a different strategy.

We use STAN to run an initial model that identifies the two most extreme decisions as anchors for the second run. For this first run, we use one single chain that is not identified with regards to rotational invariance and calculate the median of all parameters θ . For the second run—the ‘proper’ estimation—we use these two anchors to identify rotational invariance. If in a draw the right anchor ends up to the left of the left anchor, we have reason to believe that the model flipped and therefore multiply all θ parameters with -1 . In this second run, we draw overall 8’000 times from the posterior across multiple chains to effectively explore its central tendency and variation.⁸

5 Forum-Shopping and Forum-Selling in Germany’s Lower Courts

Are some lower courts systematically more plaintiff friendly than other courts? We now study decisions by German *Landgerichte* (district courts) and *Oberlandesgerichte* (regional courts of appeal) with our model to see whether we can uncover systematic evidence for forum shopping and forum selling. While it has been documented that plaintiffs engage in ‘forum shopping’—and judges in ‘forum selling’—if respective institutional incentives are in place (e.g. Bechtold, Frankenreiter and Klerman, 2019), Germany is a particularly unlikely case for this kind of behavior: Its civil law system is known

⁸Please see the appendix for convergence diagnostics.

to generate decisions in a bureaucratic way, and most importantly, without recording the individual votes of the participating judges who can be seen as apolitical career civil servants. As a matter of fact, Germany is often cited as the exemplary case of an impartial and objective civil law system (Langbein, 1985). Extending research beyond the well-known U.S. context, Bechtold, Frankenreiter and Klerman (2019) interview attorneys, judges and court officials and document mechanisms through which forum selling in fact also occurs in Germany. In areas like press law and antitrust 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. Earning a reputation for a tendency in rule making is a successful strategy when “selling” the own forum—anecdotal evidence also echoed in recent press reports (Dahlkamp and Schmid, 2014; van Lijnden, 2016). Following the lead of Bechtold, Frankenreiter and Klerman (2019), we therefore decide to study case-spaces in press law and antitrust. Our quantitative analysis will test expectations generated by their case studies on a large sample of decisions.

5.1 Collecting Data for a Common Case-Space

How to identify an appropriate set of court decisions that can be compared to one another? And how to turn citation patterns into a document-source matrix that can be fed into the measurement model? Typically, an analyst would first have to curate cases that relate to the exact same legal topic and the same regulatory context—after all each decision needs in theory be able to refer to each legal source. She then identifies and counts all legal sources the judge was citing and turns the result into a table: Each row corresponds to a decision and each column carries the count for the respective citations of a legal source. To save the effort of relying on manual labour for this task, we show how to employ technology and, at the same time, keep this process transparent and replicable. We use methods from information retrieval for collecting a well defined set of decisions, then identify the citations, and finally generate the decision-source matrix.

Our data is from the legal database *Juris*⁹ which comprises a comprehensive digital collection

⁹The data base is comparable to services like *Westlaw* for the British or the U.S. context.

of all available German court decisions. Each of *Juris*' records contains not only the complete text, but also further meta data such as titles, dates or the respective court. The data provides information on the further fate of a legal proceeding within the judicial hierarchy, all citations a court decision makes (backward citation), as well as other court decisions that refer to a certain decision (forward citation). The annotation of the corpus of decisions with meta data makes the further processing particularly easy.

Granted access to the backend of *Juris*, we worked with a data base (*MongoDB*) that was indexed with a *Lucene* based search engine (*ElasticSearch*). This infrastructure allowed us to make full use of search engine functionalities and query this corpus as we saw fit. While we collected the data for cases on press law with this infrastructure, we also used a second approach in the antitrust case for the sake of reproducibility. There, we accessed the data through the frontend, working with *Juris*' search functionality of their homepage.

First we need to define the set of decisions that belong to the common case-space. For the application to press law, we begin with selecting key words that describe the two areas we are interested in: decisions regarding claims for compensation and decisions concerned with claims for injunction.¹⁰ Acknowledging human limitations in devising dictionaries for direct document selection (e.g. Beauchamp, 2017; King, Pan and Roberts, 2013; Puglisi and Snyder Jr, 2011), we strive to mitigate this bias and cast a fairly wide set of terms to retrieve a large collection of 100 documents. A trained human coder then uses reasoning to identify those cases that truly belong to a common case-space. In result, we select nine cases on privacy infringements through the publication of photos in print media and a respective claim for compensation (dataset d_{1A}). We also collect six decisions on privacy infringement through criminal act allegation in the media and the corresponding claim for injunction (dataset d_{2A}).¹¹

In the light of these fairly small sets, we seek to reliably enlarge our samples. We take the titles of the decisions from the two already identified sets d_{1A} and d_{2A} as query terms.¹² The search engine identifies relevant decisions on the basis of the cosine similarity between the titles of the decisions

¹⁰See the appendix for the respective dictionaries.

¹¹See the appendix for the analysis of the datasets d_{1A} and d_{2A} .

¹²We use the titles instead of the full text of these decisions since the former yielded much better query results.

in the seed sets and the full text of the decisions in the data base.¹³ We find that the top 25 most similar results reliably belong to a common case-space and use those decisions as the enlarged set for claims for compensation (dataset d_{1B}) and enlarged set of decisions on privacy infringement through criminal act allegation in the media and the corresponding claim for injunction (dataset d_{2B}).¹⁴

For antitrust law we proceed in a similar way—this time, however, we access the database through the *Juris* online portal. We search with key words on antitrust¹⁵ for decisions from German district courts (*Landgerichte*). Again, a trained coder reviewed all resulting 64 decisions and clustered them into reasonable topics. We identified twenty decisions in the common case-space regarding damage claims following from a cartel’s action (dataset d_3). Given the size of this set of decisions, we decide to not further enlarge our sample.

Once the decisions in each set are defined, identifying the references is straightforward: *Juris* already enriched all decisions with meta data. When accessing the original *Juris* data for press law with the own search engine, all sources of a decision are available as a list. In turn, when querying the data via *Juris*’ homepage, downloading each decision’s HTML page allows to extract citations on the basis of embedded hyperlinks.¹⁶

As a third and final step, we count the sources in all court decisions. We define a legal source by its section (*Paragraph*) and paragraph (*Absatz*). The same law and section, but with a different paragraph counts 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); but also 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

¹³See the appendix for a more in-depth description of our approach. In the judicial politics literature, high level approaches such as the use of plagiarism software have also found attention (Hinkle, 2015). On the usage of cosine-similarity as a center piece of search engines, see Manning, Raghavan and Schütze (2009).

¹⁴The threshold of 25 cases was selected on an empirically informed basis: We found that the resulting decisions were part of the same case-space.

¹⁵See the respective dictionary in the appendix.

¹⁶On some rare occasions, decisions use the ‘Ibid.’ citation style (‘a.a.O’ in the German context) and the *Juris* data does not fully annotate this information (see also Coupette and Fleckner, 2018). The count of existing links between a decision and a legal source is potentially lower if a court uses this citation style. If anything, this bias makes our final estimates more conservative.

Court (BVerfG) as a specialized court of higher order. We also consider mentions of academic articles. Using the meta data, we then construct the respective decision-source matrix Y_{ij} that identifies all referenced legal sources (j) for each pre-selected decision (i). When counting the sources, we leave out all citations to idiosyncratic legal sources, meaning any legal source which is mentioned by only one decision.¹⁷

5.2 Application to Press Law

Let us now analyze the court decisions in press law. Jürgens (2014) gathers data on press law caseload at Germany's regional courts. He finds that between 2010 and 2012, Berlin (28.67%), Hamburg (22.32%), and Cologne (*Köln*) (11.94%) adopt the bulk of all court decisions with the rest of the caseload shared among all 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. Adding to the data from case loads, authors substantiate their claims with anecdotal evidence from either 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 rather the result of a concentration of media companies—and that includes an ecosystem of specialized press lawyers who can serve them (Dölling, 2015; Höcker and Brost, 2015). Anecdotal evidence suggests that this is the case in particular for Berlin (Dölling, 2015). Answering the call for an analysis that “meets the standards of an empirical science” (Dölling, 2015, 130, *own translation*), we investigate these claims with our measurement model. Are the courts in Hamburg and Cologne really more likely to adopt decisions that are friendly to litigants and hostile towards the press?

¹⁷These idiosyncratic legal sources do not offer any information about locations in the context of our model and leaving them out speeds-up estimation.

5.2.1 Estimates Of Decision Locations

Based on these 25 most (cosine) similar decisions each, we generate the decision-source matrices for d_{1B} and d_{2B} , omitting idiosyncratic references. Figure 2 displays the median estimates with a circle and the bars indicate the core 90% credible interval.¹⁸ In the first set of court 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. On the top of each figure we present the mean difference between the positions of decisions from Cologne and Hamburg, in contrast to all others. As the credible intervals indicate there is evidence in favor of a difference between the two groups.

In sum, we find evidence that lends itself to the experts’ and journalists’ expectations. In the context of the analyzed decisions in press law, lower courts throughout Germany show systematic differences in their judgments on comparable cases when it comes to compensation or injunction in privacy infringement cases.

5.2.2 A Qualitative Case Study to Assess the Model Validity


We now qualitatively assess the validity of our estimates and study three decisions from the extended set of cases on privacy infringement and the claim for compensation (d_{1B}): Figure 3 charts two decisions that are scaled at similar positions on the latent dimension (LG Hamburg, 324 O 161/15 and LG Cologne, 28 O 466/14) and one decision that is scaled at the opposite end (LG Munich, 9 O 23075/07).¹⁹ 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 Cologne, 28 O 466/14. In return, the probability that LG Hamburg, 324 O 161/15 and LG Cologne, 28 O 466/14 are different from one another is relatively low. If our approach is valid, then similar court decisions will derive their legal argument using similar legal sources and the court decision that is distinct should rely on different legal sources.

In the appendix we report a full legal interpretation of the arguments and sources the three courts are using. Figure 4 summarizes our findings and illustrates that our scaling approach is indeed capable

¹⁸We report the estimates for the two smaller sets d_{A1} and d_{B1} in the appendix.

¹⁹To calculate these first differences, we subtract the respective position estimates for posterior draws from one another.

Privacy Infringement—Compensation



Privacy Infringement—Injunction




Figure 2: Estimated Locations of Written Decisions ($\hat{\theta}_i$). Set of Decisions Using a ‘More-Like-This’ Query. On the Top: Mean Difference Between the Decisions from Courts in Cologne and Hamburg and All Others. Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

Privacy Infringement—Compensation First Differences of Three Decisions




Figure 3: 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.

of differentiating between different nuances in the main body of the decisions, where courts argue the cases. 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. 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 refers to different sources than the other two decisions. Ultimately, referring to different norms and case-law leads to different outcomes favoring either the litigant or the defendant. This speaks to the validity of our scaling approach: the median scores estimated for the Hamburg and Cologne decision are located at similar ends of our common case-space while the Munich decision is placed at the opposite end (Figure 2).

5.2.3 Assessment of Model Fit using Posterior Predictive Checks

Finally, we also want to assess whether the model actually fits the data generating process that we posit using posterior predictive checks. We therefore predict the outcome variable on the basis of the last 50 parameter draws and plot the predictions against the original data. Figure 5 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 chart the respective counts horizontally and show the square of their occurrence vertically.²⁰ Credible intervals are at the core 80%, 60%, 40% and 20% of the distribution. The solid lighter line depicts the median prediction across all models. The darker line plots the observed data. Our predictions mostly match the observed data and we conclude that the models do a good job in predicting the observed counts. We are thus confident that our systematic and our stochastic component correctly model the true underlying data generating citation count process.

²⁰In line with Clark and Lauderdale (2012), we use the counts' squared values for better visibility.

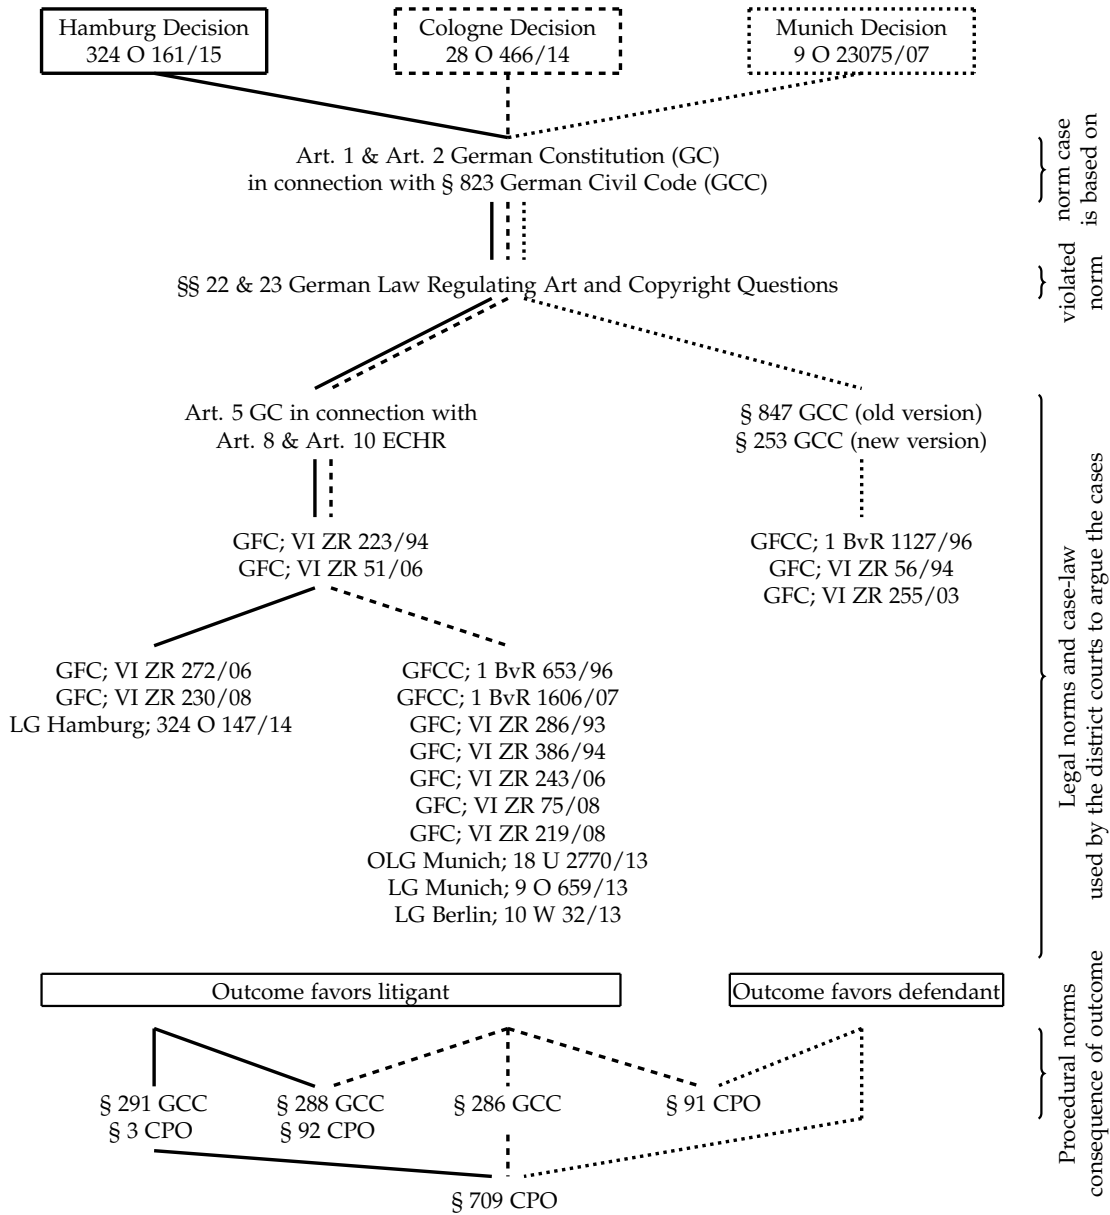



Figure 4: Summary of Legal Argumentation in Three Decisions from the Common Case-Space on Privacy Infringement and Claims for Compensation.

Privacy Infringement—Compensation



Privacy Infringement—Injunction




Figure 5: 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 in Red Are Based on the Last 50 Parameter 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.

5.3 Application to Antitrust Law

Antitrust is yet another area in German law that offers incentives for plaintiffs or courts to behave strategically. Plaintiffs can file their case either at the seat of a cartel member or at the place where mischief occurred. Often, cartels operate nationally—which means that the plaintiff can chose from the overall 24 regional courts when filing the complaint. Courts themselves are interested in attracting cases for their highly visible and exciting nature.

Bechtold, Frankenreiter and Klerman (2019) indeed find some evidence in favor of forum selling for the regional court in Mannheim. However, they are more cautious in their conclusions with regards to other cases. According to their interviews, the three regional courts in Cologne, Dortmund and Mannheim are particularly plaintiff friendly. In contrast, there are other courts such as Kiel, Leipzig, Düsseldorf or Munich that are less so.

Figure 6 plots the result for our twenty decisions. The courts in Mannheim and Dortmund tend to adopt decisions that are on one side of the spectrum. Decisions from Cologne seem to have a more

central position in our sample. Other courts either occupy the centre or display estimated locations of decisions on the centre-right. The small plot on top charts the difference between the decisions from courts in Cologne, Dortmund and Mannheim versus all other courts. The results support the existing qualitative evidence. Our results show that the mean differences between decisions of those three courts and the other courts for the common case-space in antitrust is systematic.

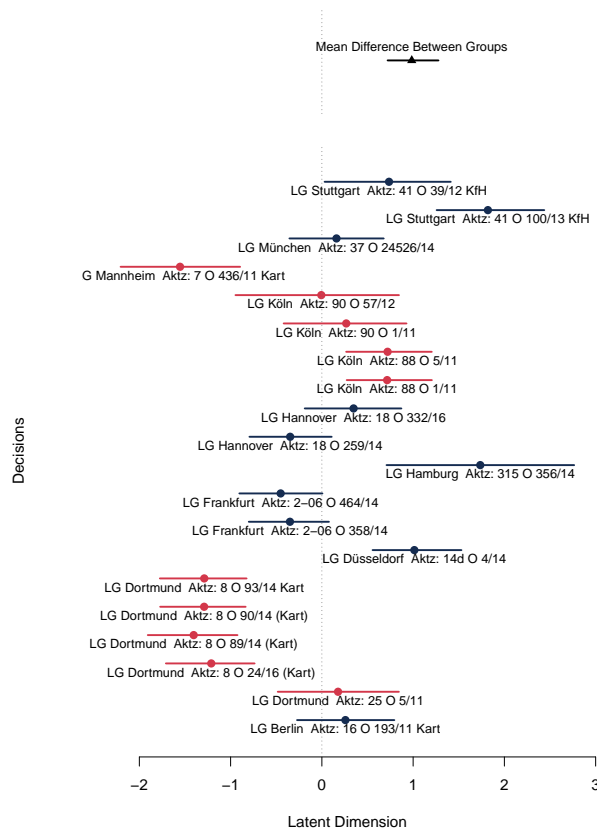


Figure 6: Estimated Locations ($\hat{\theta}_i$) of Written Decisions on Antitrust. On the Top: Mean Difference Between the Decisions from Courts in Cologne, Dortmund and Mannheim and All Others. Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

6 Conclusion

Legal scholars and political scientists are not only interested in understanding the decisions judges make, but also seek to analyze the policy implications of written decisions. Previous strategies usually

rely on voting data of the involved judges (e.g. Clark and Lauderdale, 2010). But in cases where no such records exists, empirical strategies to scale decisions in a common space are still missing. We show how to estimate and consequently compare positions of various actors in a common case-space of court decisions using citations networks. Our model rests on the notion that ‘similar’ decisions, i.e. decisions with similar locations in a common space, express their similar legal reasoning through similar citation patterns.²¹

To showcase the abilities of our model, we study a particularly challenging case: the German legal system. Qualitative evidence indicates that German lower courts have a systematic bias in some legal areas, engaging in *forum selling* in press law and antitrust (Klerman and Reilly, 2016; Bechtold, Frankenreiter and Klerman, 2019). We corroborate these findings. Indeed, there is evidence that some lower courts—in press law Cologne and Hamburg, and for antitrust Cologne, Dortmund and Mannheim—do take systematically biased decisions. We can also replicate our findings for different sets of decisions. In addition, we show how to find appropriate decisions in a large legal corpus that all belong to a common case-space, thus leading the way to analyzing larger sets of court decisions without human intervention.

Of course, there are limits to what our model can do. Just like any other item-response theory approach—for example when locating parliamentarians on a left-right spectrum—all units that are being scaled need to reside in the same latent space. For our legal application, this means that all decisions have to concern the same legal matter and this also includes a constant regulatory contexts: Any major legislative change that affects the legal sources available to the judge is likely to introduce bias. Future research might build on our model to automate the process of identifying decisions that belong to a common case-space even further. Moreover, although our model makes moderate assumptions about how courts refer to legal sources to justify their legal arguments, the evidence we provide comes only from one country, Germany, which is a typical civil law country. Future research might be able to provide more evidence that this model is useful to locate and compare court decisions

²¹To a degree, our understanding of references blur the conceptual difference between courts applying the law—in our terminology referring to legal provisions—in contrast to courts providing arguments—here referring to former decisions or literature. We leave it to further research to develop models that make full use of this conceptual distinction in their empirical strategy.

within an appropriate case-space in other contexts as well.

Our model enriches the toolbox for an interdisciplinary group of scholars who study judges and their decisions quantitatively. Legal scholars might be interested in unveiling the development of doctrine across time and space. Political scientists might want to understand the nexus between institutions and power. Others, like economists or sociologists, might rather be concerned with the effects of court decisions on business or societies at large. Common space positions of court decisions are the foundation for the development and testing of powerful analytical models and may find their application in any of these fields.

Lastly, our insights have the potential to be useful beyond academia. Locating written lower court decisions in a common space has also important practical implications for law firms. 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 courts 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. Thus, locating actual decisions in a common case-space is very helpful for all who want to quickly spot similar or opposing decisions. With our model they can retrieve an easy roadmap capable of guiding them in their professional effort.

References

- Alschner, Wolfgang and Damien Charlotin. 2018. “The growing complexity of the International Court of Justice’s self-citation network.” *European Journal of International Law* 29(1):83–112.
- Barberá, Pablo. 2015. “Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data.” *Political Analysis* 23(1):76–91.
- Beauchamp, Nicholas. 2017. “Predicting and Interpolating State-Level Polls Using Twitter Textual Data.” *American Journal of Political Science* 61(2):490–503.
- Bechtold, Stefan, Jens Frankenreiter and Daniel Klerman. 2019. “Forum selling abroad.” *Southern California Law Review* 92(3):487–556.
- 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.
- Cameron, Charles M. and Lewis A. Kornhauser. 2017a. “Chapter 2: What Courts Do . . . And How to Model It.” *Working Paper* (June).
- Cameron, Charles M. and Lewis A. Kornhauser. 2017b. “Chapter 3: What Do Judges Want? How to Model Judicial Preferences.” *Working Paper* .
- 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.
- Charlotin, Damien. 2017. “The Place of Investment Awards and WTO Decisions in International Law: A Citation Analysis.” *Journal of International Economic Law* 20(2):279–299.
- Choi, Stephen J. and G. Mitu Gulati. 2008. “Bias in Judicial Citations : A Window into the Behavior of Judges ?” *Journal of Legal Studies* 37(1):87–129.
- 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 L. 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 Multilevel/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.
- Hanretty, Chris. 2012a. "Dissent in Iberia: The ideal points of justices on the Spanish and Portuguese Constitutional Tribunals." *European Journal of Political Research* 51(5):671–692.
- Hanretty, Chris. 2012b. "The Decisions and Ideal Points of British Law Lords." *British Journal of Political Science* 43(3):1–3.
- Hinkle, Rachael K. 2015. "Into the Words: Using Statutory Text to Explore the Impact of Federal Courts on State Policy Diffusion." *American Journal of Political Science* 59(4):1002–1021.
- Höch, Dominik. 2018. "Alles neu, alles anders? Das BVerfG nimmt sich der Gerichtspraxis im Presserecht an."
URL: 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. 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.

- Jakab, András, Arthur Deyevre and Giulio Itzcovich, eds. 2017. *Comparative Constitutional Reasoning*. 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.
- King, Gary, Jennifer Pan and Margaret E Roberts. 2013. “How censorship in China Allows Government Criticism But Silences Collective Expression.” *American Political Science Review* 107(2):326–343.
- Klerman, Daniel and Greg Reilly. 2016. “Forum Selling.” *Southern California Law Review* 89(2):241–316.
- Kompa, Markus. 2012. “Experiment zum fliegenden Gerichtsstand.”
URL: <https://heise.de/-3394265>
- 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.
- 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.
- Lupu, Yonatan and Erik Voeten. 2012. “Precedent in International Courts: A Network Analysis of Case Citations by the European Court of Human Rights.” *British Journal of Political Science* 42(02):413–439.
- Lupu, Yonatan and James H. Fowler. 2013. “Strategic Citations to Precedent on the U.S. Supreme Court.” *Journal of Legal Studies* 42(1):151–186.
- 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.

- 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* .
- Owens, Ryan J. and Justin P. Wedeking. 2011. "Justices and legal clarity: Analyzing the complexity of U.S. Supreme Court opinions." *Law and Society Review* 45(4):1027–1061.
- Owens, Ryan J. and Justin P. Wedeking. 2012. "Predicting Drift on Politically Insulated Institutions: A Study of Ideological Drift on the United States Supreme Court." *The Journal of Politics* 74(2):487–500.
- Petersen, Niels and Emanuel V. Towfigh. 2017. "Network Analysis and Legal Scholarship." *German Law Journal* 18(3):695–700.
- Poole, Keith T. and Howard Rosenthal. 2007. *Ideology and Congress*. Transaction Publishers.
- Posner, Richard A. 2000. "An Economic Analysis of the Use of Citations in the Law." *American Law and Economics Review* 2(2):381–406.
- Puglisi, Riccardo and James M Snyder Jr. 2011. "Newspaper Coverage of Political Scandals." *The Journal of Politics* 73(3):931–950.
- 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):12.
- 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.

Scaling Court Decisions with Citation

Networks

Appendix

A Proof of Concept for the Model

We devise toy examples to show that our model indeed correctly picks up positions in a way we expect. Let us begin with a baseline experiment. The artificial data we constructed in Table A.1 specifies whether a court decision cites a certain source or not. There are five decisions that refer to five different legal sources. If the decision refers to a legal source, it carries a 1, otherwise it has a 0. The way the data is set up, it is straightforward to see that all decisions should be distributed symmetrically and at equal distances in the case space.

Table A.1: Toy Data 1: Does a Decision Refer to a Source?

	Source 1	Source 2	Source 3	Source 4	Source 5
Decision 1	1	1	0	0	0
Decision 2	0	1	1	0	0
Decision 3	0	1	1	1	0
Decision 4	0	0	1	1	0
Decision 5	0	0	0	1	1

We compare these results with a null-model where we randomly rewire the citation matrix. Figure A.1 depicts the connections—on the left the systematic data and on the right the randomly rewired matrix. We then measure the ideal points using the model from the main paper—the only difference is to select the appropriate link function for this kind of data, i.e. a logit link function. Overall, we take 8’000 draws from the posterior. The experiments are encouraging: Figure A.2 shows on the left panel for the systematic data that as expected all estimates are distributed with equal distance in the case space. Credible intervals are also well behaved. The estimated locations we




Figure A.1: Citation Source Matrices. Connections in Dark Blue.

can retrieve from the randomly rewired citation matrix are—as expected—quite random. Figure A.2 indicates on the right the results with random ideal points and overlapping credible intervals.

Recording only whether a court refers to a legal source or not is quite a strong assumption. It seems much more realistic to also think about how often a court is citing a source. While a decision considers a dissenting legal source, it might refer to it only once or twice. But a legal source that is relevant will be referred to much more often. We therefore also estimate ideal points on the basis of the more realistic data structure in A.2.

Table A.2: Toy Data 2: How Often Does a Decision Refer to a Source?

	Source 1	Source 2	Source 3	Source 4	Source 5
Decision 1	10	5	1	1	1
Decision 2	1	10	5	1	1
Decision 3	1	4	7	4	1
Decision 4	1	1	5	10	1
Decision 5	1	1	1	5	10

Again, we also rewire this count data and randomly scramble the citation counts. Figure A.3 shows again both resulting citation matrices. On the left the systematic data and on the right the

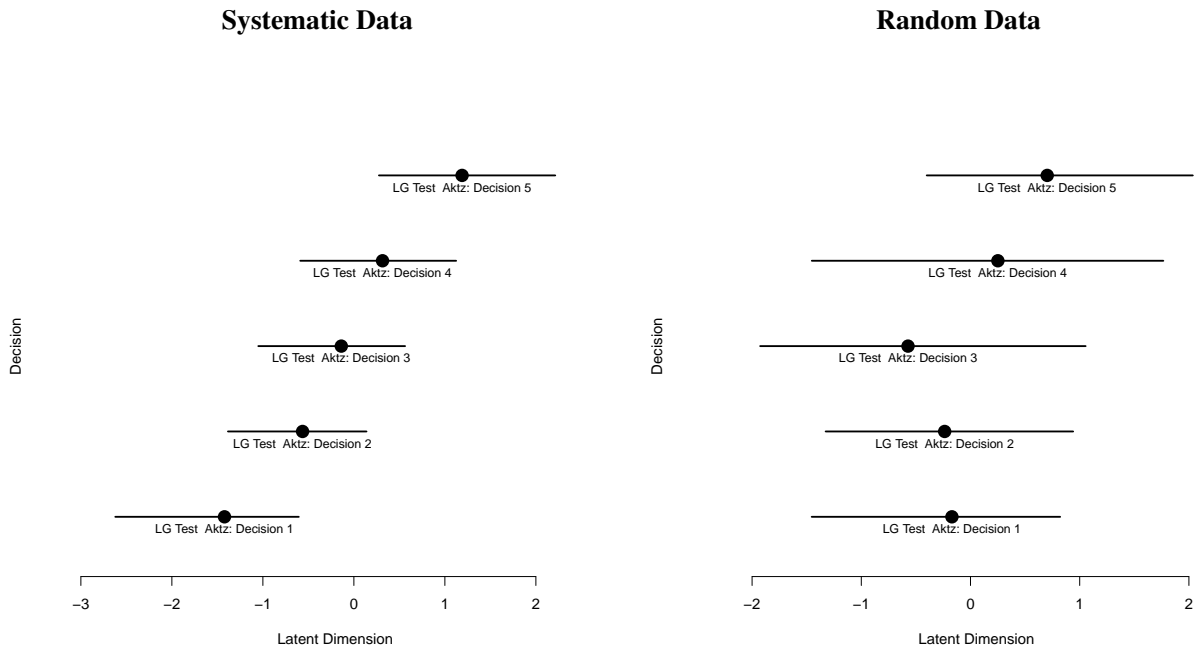


Figure A.2: Estimating the Position of Decisions on the Basis of the Data in Table A.1 and the Rewired Data.

random permutation. All decisions cite all sources—however they do so to a different degree. The resulting citation graph with weighed edges can not be estimated with a logit link function. Count models such as the poisson link function allows to appropriately account for the data generating process.

We use the exact same setup to estimate the bayesian Model—including also the flipping to solve rotational invariance—and sample overall 8’000 draws from the posterior. Results in figure A.4 show that the model retrieves ideal points that reflects the data. On the left, we clearly see the systematic pattern from the citation counts also emerging in the positions of the decisions. Credible intervals indicate that the model is capable to handle the count input data from table A.2 well. In contrast, the model on the right is an image of the random data. The estimated locations are again without a clear pattern and the credible intervals are wide.




Figure A.3: Citation Source Matrices. Connections in Blue. Darker Shadings Represent a Higher Count.




Figure A.4: Estimating the Position of Decisions on the Basis of the Data in Table A.1 and the Rewired Data.

B Enhancing the Query

The first step in data collection is to identify a suitable set of decisions that belong to the same legal case-space.

B.1 Dictionaries for Querying the Data Base

For press law, we are accessing Juris data directly with a *Lucene* based search engine (*ElasticSearch*). We define the following two dictionaries for the initial seed query. Table B.1 relates to claims for compensation. In a similar vein, Table B.2 concerns claims that demand injunction.

Table B.1: Dictionary to Query the Data Base for Decisions on Privacy Infringement Claiming Compensation.

Presserecht, Presse, Pressefreiheit, Presseerzeugnis, Äußerung, Interview, Darstellung, Medien, Meinungsfreiheit, Meinung, Meinungsäußerung, Persönlichkeitsrechtsverletzung, Persönlichkeitsrecht, Schadensersatz, Schadensersatzanspruch, Schadensersatzberechnung, Schadensersatzklage, Schadensersatz, Schadensfeststellung, Schadenshöhe

Table B.2: Dictionary to Query the Data Base for Decisions on Privacy Infringement Demanding Injunction.

Presserecht, Presse, Pressefreiheit, Presseerzeugnis, Äußerung, Interview, Darstellung, Medien, Meinungsfreiheit, Meinung, Meinungsäußerung, Persönlichkeitsrechtsverletzung, Persönlichkeitsrecht, Unterlassung, Unterlassungsanspruch, Unterlassungsklage, Unterlassungsverfügung, Unterlassungserklärung, Unterlassungsangebot, Unterlassungspflicht, Unterlassungsantrag, Unterlassungsverpflichtung

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 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*”.

B.2 Dictionaries for Querying Juris on their Homepage

To analyze cartel cases, we query the Juris homepage directly. The two terms that make up the dictionary are *Kartellrecht* and *Schadensersatz*. The dictionary is much shorter than for press law, because we have to abide by the more limited query functionality of the *Juris* frontend. We also restrict the search to lie between 01.01.2012-01.01.2018, with the author being a district court (*Landgericht*).

B.3 Calculating the Similarity between Query Terms and Documents

For our four analyses on press law, we query the data base with the search engine. We first use a seed list with terms we are interested in (Table B.1 and Table B.2). But to expand our sample, we also query the data base with the titles of the decisions that are in these two sets.

How does the search engine evaluate the similarity between the input we provide and the corpus in the data base? We use the cosine similarity as the key quantity to measure the similarity between a query phrase and a document in the data base. While in the legal literature, more high level approaches such as the use of plagiarism software have been applied for the same task (Hinkle, 2015), we rely on the cosine similarity, since it is a fundamental measurement that finds widespread application in many search engines.¹ When looking for similar legal documents in the data base, the algorithm considers the complete text of all decisions. 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 that contains the whole vocabulary of the corpus on the first dimension. The second dimension holds the word counts of each document in the corpus.² 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

¹For a more in-depth treatment, see Manning, Raghavan and Schütze (2009).

²In 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.

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.³ Implementing this measure, we concatenate the titles from the query set to one single query document d and find the most similar court 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.

C Decision Source Matrix

We also provide a quick overview over the decision-source matrices we find and decide to display them visually.

C.1 Press Law: Hand Selected

Figure C.1 provides an overview over the resulting decision-source matrices Y_{ij} for the sets d_{1A} and d_{2A} . Each row represents a court decision, each column a legal source. The darker the color, the more often a decision refers to a particular legal source. For the cases on compensation, the figure on the left displays a number of long vertical lines that indicate a high degree of overlap: These decisions refer to similar legal sources—but some more and others less often. In addition, there are a number of legal sources that are being referenced by only a few court decisions. The decision-source 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 court decisions, only.

³The *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 (last accessed April 2021).

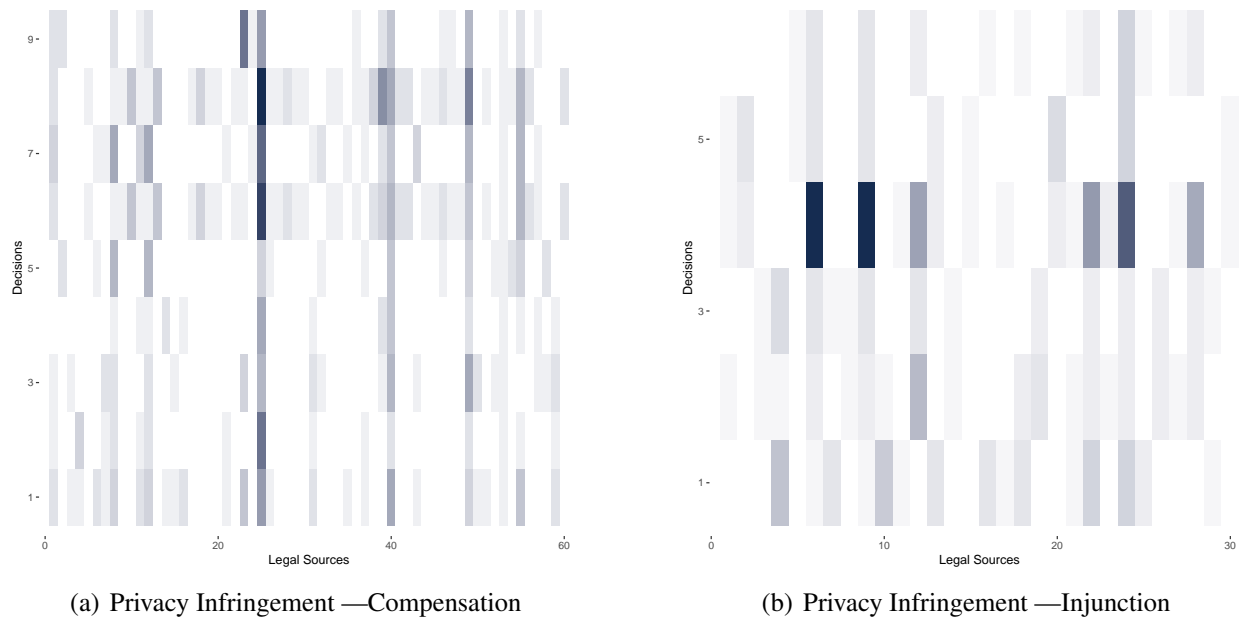


Figure C.1: Exact Query. Coding of the Decision-Source Matrix on the Basis of Metadata. The Darker the Shading, the More Often a Court Decision Refers to a Legal Source.

C.2 Press Law: Hand Selected and Query Expansion

Figure C.2 displays the data from the sets d_{1B} and d_{2B} . For the decisions on compensation, the matrix is quite well-behaved and shows a substantial degree of overlap between those decisions and the legal sources. We can clearly identify five legal sources that are being referred to by a large number of decisions. In addition, there are legal sources that are mentioned by some decisions, which in the end provide most of the analytic leverage. The decision-source matrix of the second set of written decisions on injunction shows less overlap. Only two legal sources are apparently widely mentioned.

C.3 Anti Trust: Hand Coded

We also chart the decision-source matrix for the set of decisions in antitrust in Figure C.3. This set has considerably more legal sources than the other sets in press law. Clearly, there seems to be a core doctrine that courts typically refer to. It is easy to identify it on the left of the figure in darker shading.

Due to the manual annotation process of the legal sources displayed in the left panel, the sources are ordered in a way that generates a triangle. Indeed there could be the impression of a systematic pattern due to time dependency. However, the figure on the right hand puts this impression into




Figure C.2: 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 Court Decision Refers to a Legal Source.

perspective. It displays the same citation data with one important difference: the sources are ordered by year. If there were indeed a systematic time dependency, we would expect to see the same triangle we observe on the left, which is not the case. Instead, the pattern is apparently a consequence of the manual coding.

D Estimating Decision Locations

D.1 Estimating the Location of Further Decisions

Here we present the estimated locations for the two sets of decisions d_{A1} and d_{B1} we collected with a key search query. The resulting position estimates are in line with *ex-ante* expectations based on expert knowledge and media reports. Figure C.4 depicts the estimated decision locations. The figure displays all decision locations in our samples from Cologne 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 plots on top sum-

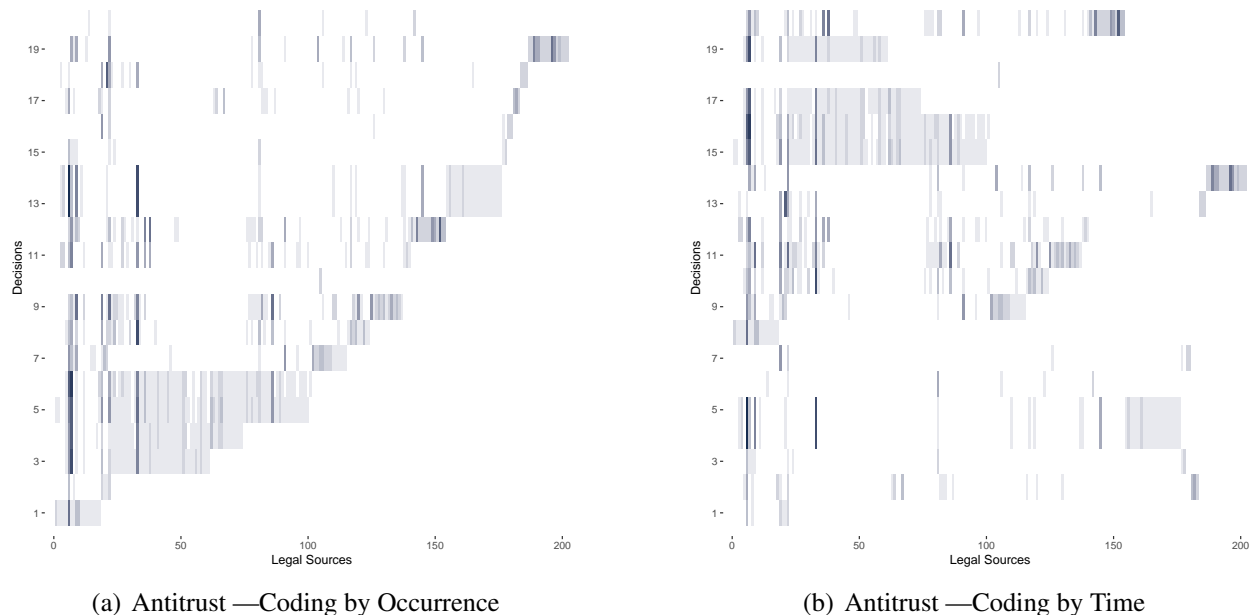


Figure C.3: Set of Decisions Accessed via *Juris* Homepage. Coding of the Decision-Source Matrix on the Basis of Metadata. The Darker the Shading, the More Often a Court Decision Refers to a Legal Source.

marize the mean difference between the decisions from the courts in Cologne and Hamburg *vis-à-vis* all other courts. For the decisions related to compensation in the panel on the left, we observe two outliers—both from Cologne. However, even though Hamburg and Berlin are also known as friendly towards compensation, the decisions do not show systematic differences. In the decisions related to injunction on the right panel, the decisions from Cologne and from Hamburg cluster, i.e. they are mapped onto a similar location in the case-space. Heidelberg, the only other court in the sample, is distinctly situated on the right. The estimated locations are in line with anecdotal evidence from media reports and from experts.

D.2 Estimating the Location of Sources

Our model also estimates the location of the legal sources within the same case-space. As an illustration, Figure D.1 shows the estimated locations ($\hat{\phi}_i$) of every legal source for the two sets of decision in press law without query extension (d_{A1} and d_{B1}). The other sets of decisions contain even more legal sources which would make them difficult to display at the detailed level.

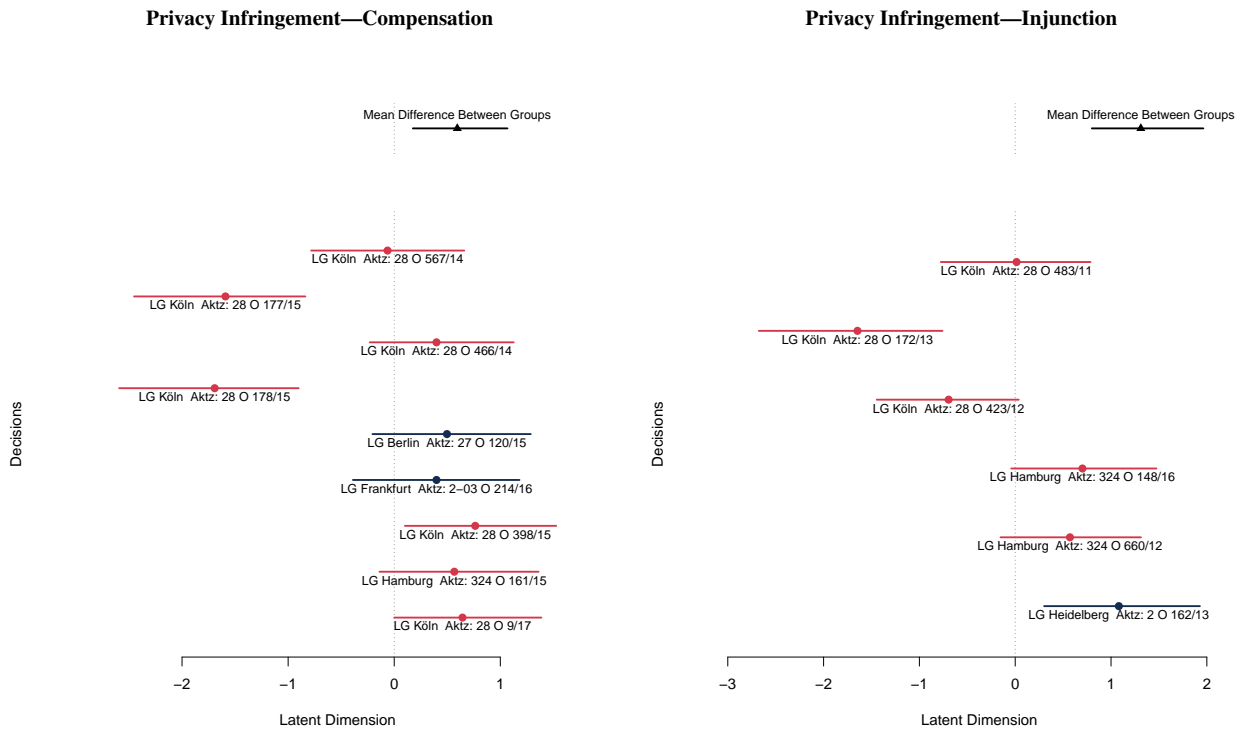
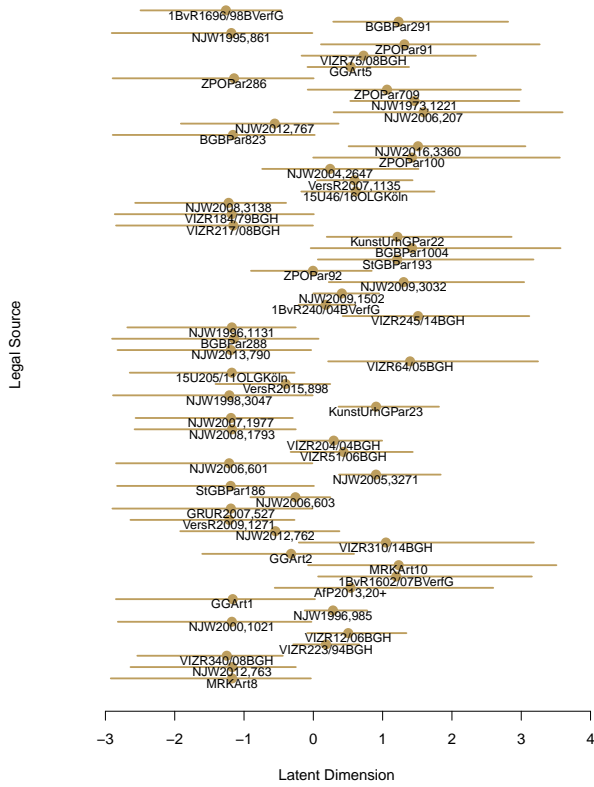


Figure C.4: Estimated Locations of Written Decisions ($\hat{\theta}_i$). Pre-selected Set of Decisions. On the Top: Mean Difference Between the Decisions from Courts in Cologne and Hamburg and All Others. Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

Privacy Infringement — Compensation



Privacy Infringement — Injunction




Figure D.1: Estimated Locations of Legal Sources ($\hat{\phi}_j$). Sets of Decisions Without Query Extension (d_{1A}, d_{1B}). Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

We do not only scale decisions but also legal sources in the same case-space. This can facilitate more fine-grained substantive interpretations of the legal argumentation that is developed within those decisions because the type and count of the legal sources provide additional information that has not been leveraged before.

E Convergence Diagnostics

We also add convergence diagnostics for the parameter θ_i in each model.

Table E.1: Convergence Diagnostics for estimated Locations. Press Law Images Exact (Case d_{A1}).

	mean	se	n_{eff}	\hat{R}
theta[1]	0.67	0.02	507.94	1.02
theta[2]	0.58	0.01	1,242.48	1.01
theta[3]	0.78	0.02	433.20	1.02
theta[4]	0.40	0.01	1,138.06	1.01
theta[5]	0.52	0.01	1,063.25	1.01
theta[6]	-1.71	0.03	367.08	1.03
theta[7]	0.41	0.01	961.74	1.01
theta[8]	-1.61	0.03	381.14	1.02
theta[9]	-0.06	0.01	1,667.67	1.00

Table E.2: Convergence Diagnostics for estimated Locations. Press Law Images MLT (Case d_{A2}).

	mean	se	n_{eff}	\hat{R}
theta[1]	-0.38	0.02	2,025.77	1.00
theta[2]	0.92	0.01	1,230.50	1.00
theta[3]	0.34	0.01	2,292.68	1.00
theta[4]	0.70	0.02	2,010.14	1.00
theta[5]	-1.13	0.01	2,059.62	1.00
theta[6]	-1.64	0.02	590.65	1.02
theta[7]	1.13	0.01	1,223.80	1.01
theta[8]	-0.55	0.01	1,365.37	1.01
theta[9]	-0.03	0.01	2,712.72	1.00
theta[10]	-0.17	0.01	3,054.80	1.00
theta[11]	0.34	0.01	3,397.89	1.00
theta[12]	1.46	0.01	974.77	1.01
theta[13]	-1.03	0.01	889.72	1.01
theta[14]	-0.28	0.02	957.51	1.01
theta[15]	-0.87	0.01	2,243.65	1.00
theta[16]	0.08	0.01	1,279.18	1.00
theta[17]	0.49	0.01	1,201.28	1.00
theta[18]	0.81	0.01	1,205.08	1.00
theta[19]	0.72	0.01	2,868.81	1.00
theta[20]	0.25	0.01	2,926.81	1.00
theta[21]	-0.30	0.01	2,524.90	1.00
theta[22]	-1.34	0.01	1,035.63	1.01
theta[23]	1.43	0.01	982.38	1.02
theta[24]	-1.05	0.01	971.21	1.01

Table E.3: Convergence Diagnostics for estimated Locations. Press Law Online Linking Exact (Case d_{B1}).

	mean	se	n_{eff}	\hat{R}
theta[1]	-1.10	0.02	1,088.12	1.00
theta[2]	-0.58	0.01	1,270.93	1.00
theta[3]	-0.71	0.01	1,293.31	1.00
theta[4]	0.70	0.01	1,200.63	1.00
theta[5]	1.68	0.02	883.28	1.00
theta[6]	-0.01	0.01	1,544.79	1.00

Table E.4: Convergence Diagnostics for estimated Locations. Press Law Online Linking MLT (Case d_{B2}).

	mean	se	n_{eff}	\hat{R}
theta[1]	0.13	0.01	1,522.14	1.00
theta[2]	0.25	0.01	2,041.67	1.00
theta[3]	0.29	0.01	2,028.76	1.00
theta[4]	-1.47	0.01	624.44	1.01
theta[5]	-0.01	0.01	2,466.49	1.00
theta[6]	0.32	0.01	1,377.14	1.00
theta[7]	0.01	0.01	1,933.98	1.00
theta[8]	-1.21	0.01	637.97	1.01
theta[9]	-0.06	0.01	1,611.72	1.00
theta[10]	1.74	0.01	635.35	1.00
theta[11]	0.28	0.01	2,434.83	1.00
theta[12]	-1.48	0.01	673.25	1.02
theta[13]	-0.20	0.01	2,742.61	1.00
theta[14]	0.18	0.01	1,661.74	1.00
theta[15]	-0.47	0.01	1,520.75	1.00
theta[16]	-0.81	0.01	754.06	1.01
theta[17]	-0.20	0.01	2,261.18	1.00
theta[18]	-0.08	0.01	1,288.60	1.00
theta[19]	2.06	0.02	603.57	1.00
theta[20]	-0.26	0.01	1,531.68	1.01
theta[21]	0.25	0.01	2,224.97	1.00
theta[22]	-0.18	0.01	1,379.84	1.01
theta[23]	2.04	0.01	782.39	1.00
theta[24]	-0.90	0.01	1,127.09	1.01

Table E.5: Convergence Diagnostics for estimated Locations. Antitrust (Case d_3).

	mean	se	n_{eff}	\hat{R}
theta[1]	-0.26	0.01	1,067.15	1.00
theta[2]	-0.18	0.01	1,399.32	1.00
theta[3]	1.22	0.01	391.34	1.01
theta[4]	1.42	0.02	343.43	1.01
theta[5]	1.30	0.02	349.32	1.01
theta[6]	1.30	0.02	349.70	1.01
theta[7]	-1.02	0.01	555.68	1.02
theta[8]	0.36	0.01	682.89	1.01
theta[9]	0.45	0.01	599.33	1.01
theta[10]	-1.74	0.01	2,153.27	1.01
theta[11]	0.35	0.01	717.60	1.01
theta[12]	-0.35	0.01	729.90	1.01
theta[13]	-0.72	0.01	537.23	1.01
theta[14]	-0.72	0.01	526.30	1.01
theta[15]	-0.27	0.01	1,463.24	1.00
theta[16]	0.02	0.02	1,114.47	1.01
theta[17]	1.56	0.02	466.83	1.01
theta[18]	-0.16	0.01	952.86	1.00
theta[19]	-1.84	0.02	255.54	1.03
theta[20]	-0.74	0.01	1,241.91	1.01

F Extended Qualitative Case Study to Assess Model Validity

We provide an in-depth analysis of the data to assess the validity of our measurement model. In this section of the Appendix, we have the space to fully consider the legal reasoning in all three cases. The median estimate of the Hamburg decision (LG Hamburg, 324 O 161/15) is to the left of the case-space in Figure 2. 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 defendant from publishing pictures before referring to the *Landgericht* (324 O 161/15, Mn 6).⁴ 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 it is 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

⁴We refer to sections of interest in court decisions using margin numbers (Mn) as shown in the *Juris* database.

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.

If our scaling approach is valid, then the citation pattern of the district court of Cologne in LG Köln, 28 O 466/14 should be similar. The median position of the Cologne decision is to left of the case-space in Figure 2, and 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. The litigant is also 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 reading 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 weighs—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) on the basis of the European Convention on Human Rights and the German Constitution (28 O 466/14, Mn 20-23).

To 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, the Hamburg decision and the Cologne decision are scaled at similar ends of the case-space in Figure 2. The general case characteristics and the legal outcomes are similar in both cases. In addition, the legal norms and the case-law used to argue the cases widely overlap.

The median position of the *Munich decision* (LG Munich, 9 O 23075/07) is scaled at the opposite end of the case-space in Figure 2 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—the regulations referred to in the Hamburg and Cologne decision. Nevertheless, in the Munich decision the judges rather cite 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. Our model finds 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 2. 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.

G Assessment of Validity Based on Decision Outcomes

If our approach is valid then the following should hold true: The estimated position of a decision ($\hat{\theta}_i$) published in one of the legal areas under scrutiny is a function of the legal sources mentioned in the decision. If the estimated position between two decisions published in the same case-space is wide then this implies that the variation in the legal source is larger compared to when the distance between the two decisions is narrow. Moreover, it is plausible to assume that the variation in citation should reflect on the substantive outcome of decisions. For example, in the application on antitrust law the LG Dortmund made a total of five decisions; four are positioned to the left relative to the fifth decision position to the right (see figure 6). Reviewing the substantive outcomes of the scaled decisions we find that in the four decisions to the left the smaller company is favored over the cartel, compared to the single decision to the right which favors the cartel.

Thus, to test the validity of our approach we manually coded the substantive outcomes of decisions by all courts that published at least three decisions in a respective case-space. If our approach is valid we should find that in one case-space the substantive outcomes of decisions show an ordered pattern. In other words, until a certain point courts take decisions that favor an individual person (in press law) or a small company (in antitrust law) and afterwards courts would take decisions that favor the press or the cartel. Figure G.1 illustrates that we find this pattern accounting for the credible intervals in all but one case.⁵ Decisions that favor the small company (red circles) in the case-space summarizing decisions on antitrust are positioned to the left relative to decisions that favor the cartel (blue circles). Comparing the red circles to the blue circles in the case of the LG Dortmund we estimated first differences which are not significant on conventional levels. In short, clusters with a mix between red and blue circles do not systematically break with the expected pattern. This is true for all courts and across all shown case-spaces, except for the cluster of decisions made by the KG Berlin

⁵We also calculate first differences which corroborate the impressions from Figure figure:validityOutcome. For the sake of the presentation we focus on the current visualisation.

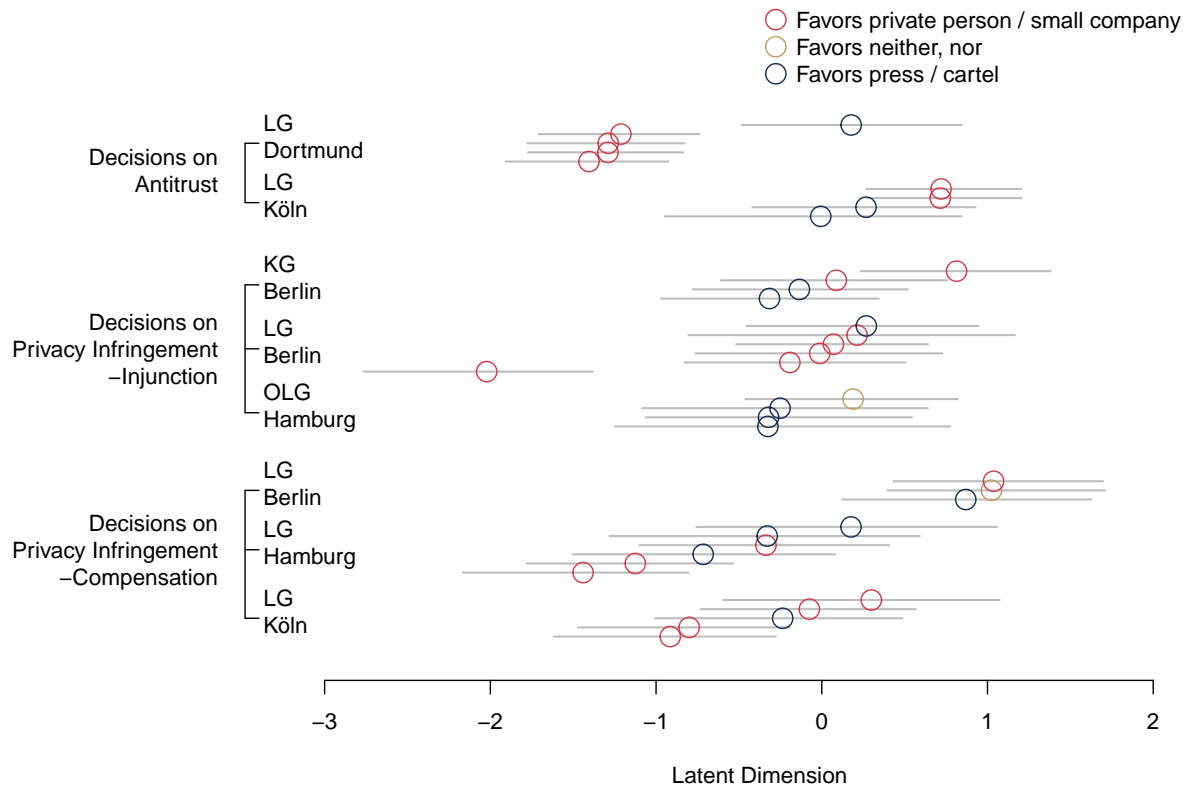


Figure G.1: Estimated Positions of Decisions and their Outcomes. Bars Represent 95% Credible Intervals


in the case-space on injunctions. Here we find systematic differences between the red circle to the right of the case-space when compared to the blue circles positioned relative to the left. In sum, with the exception of one court Figure G.1 provides robust evidence that our scaling approach produces validate locations of written decisions that at the same time predict the decision outcome. The variation in locations can also be seen when considering the variation in substantive decision outcomes accounting for estimation uncertainty.

H Null Models for Forum Shopping


Similar to the toy model in Section A of the appendix, we also generated null models for the case spaces on forum shopping. We again resample each citation-count matrix. While this keeps the digits of the decision-source matrix the same, it randomly changes their position.

Figure H.1 depicts the resulting locations. On first sight, the model seems to generate well

Privacy Infringement—Compensation



Privacy Infringement—Injunction



Antitrust





Figure H.1: Positions of Decisions from Randomly Resampled Decision-Source Matrices.


behaved estimates. However, a closer consideration of the decisions reveals that the results do not make intuitive sense. The same courts adopt decisions that are widely apart from one another—which not only goes against our theoretical expectations, but also against what we know about these decisions when reading them. For example, in the case-space on antitrust the decisions from LG Dortmund (8 O 93/14, 8 O 90/14, 8 O 89/14, 8 O 24/16) are all in the same spirit and should lead to similar locations. The locations in Figure 2 and Figure 6 reflect our knowledge about the decisions much better.

As another means of evaluating the robustness of the null models, we also add posterior predictive checks for these models in Figure H.2. Note that the results clearly indicate that the model is not a great fit for any of the three case-spaces—in particular when compared to Figure 5.

Privacy Infringement—Compensation



Privacy Infringement—Injunction



Antitrust




Figure H.2: Posterior Predictive Checks from Randomly Resampled Decision-Source Matrices.