Scaling Court Decisions with Citation Networks*

CHRISTIAN ARNOLD†
Cardiff University

BENJAMIN G. ENGST‡
University of Konstanz

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. Measurement models that support analysts in this manual labor usually rely on judges' voting records. Since these data are often not available, we instead propose a latent-variable model that uses the widely available references in court decisions to measure the decisions' latent position in their common case-space. We showcase our model in the context of forum-shopping and forum-selling of Germany's lower courts.

*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 code are available at XXX.

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

‡PostDoc, Department of Politics and Public Administration, University of Konstanz, Universitätsstraße 10, D-78464, Konstanz, Germany (benjamin.engst@uni-konstanz.de).

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

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

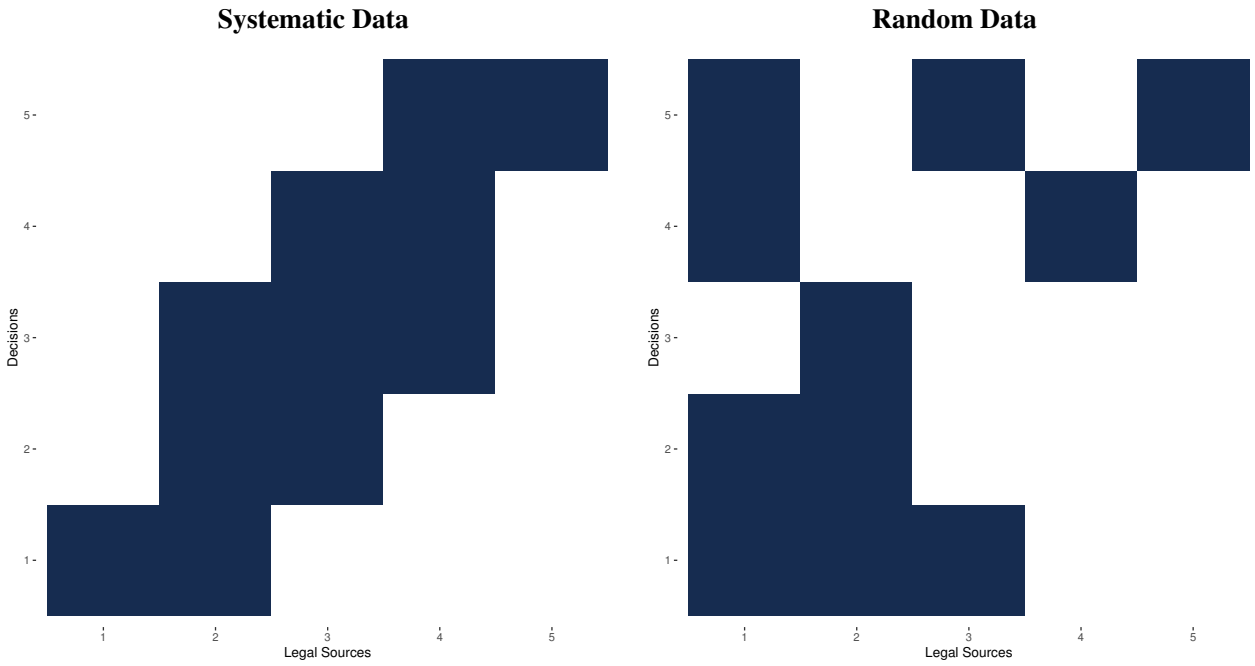

Figure A.1: Citation Source Matrices. Connections in Dark Blue.

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

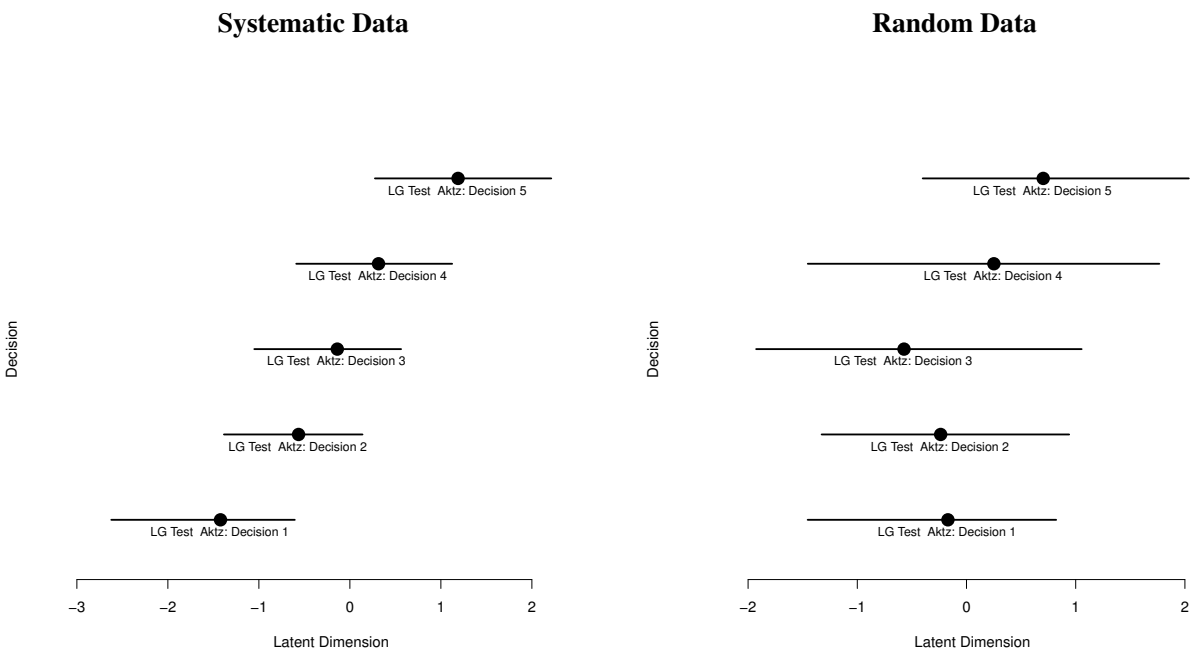

Figure A.2: Estimating the Position of Decisions on the Basis of the Data in Table A.1 and the Rewired Data.

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.

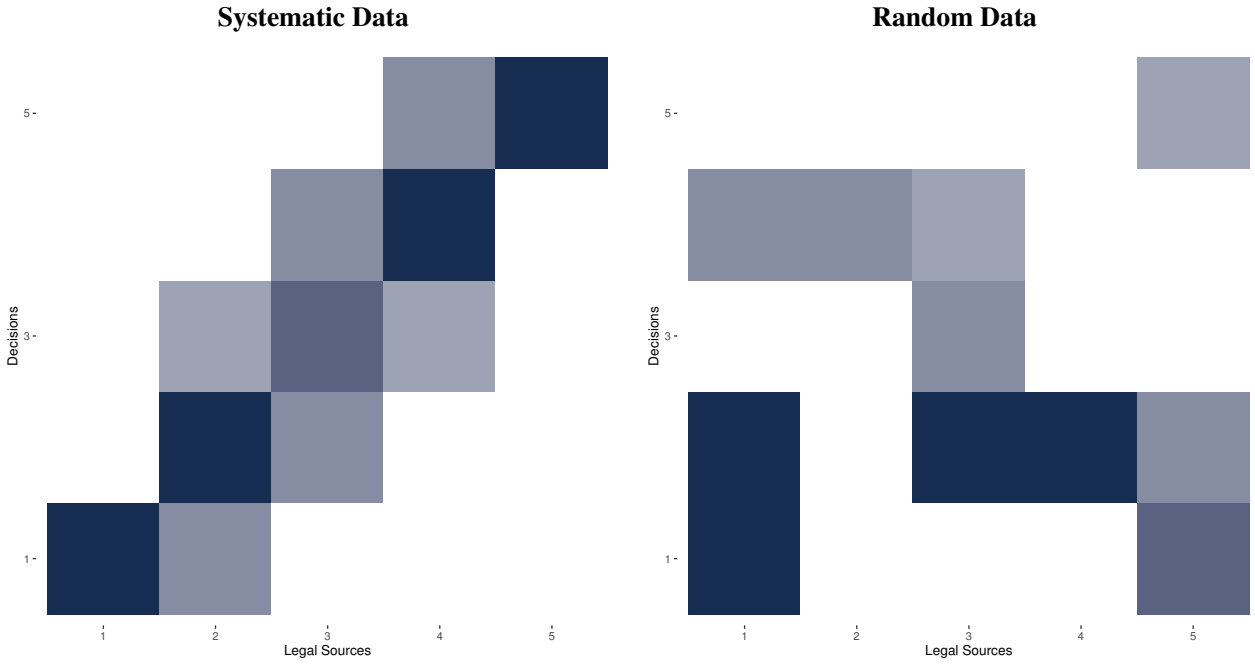

Figure A.3: Citation Source Matrices. Connections in Blue. Darker Shadings Represent a Higher Count.

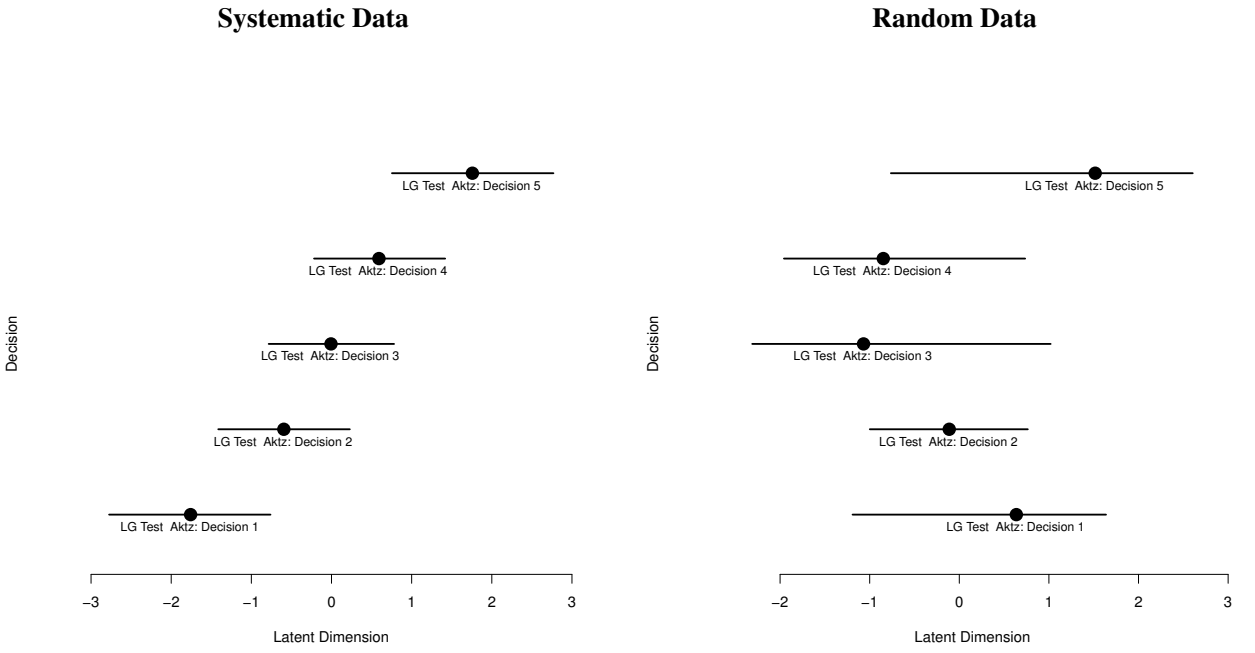

Figure A.4: Estimating the Position of Decisions on the Basis of the Data in Table A.1 and the Rewired Data.

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

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

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

cosine similarity is calculated as

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

with the enumerator 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.

[3] The *Lucene* scoring built into the *ElasticSearch* search engine further refines this measure and allows for more fine grained specifications of search queries. However, these scores are not relevant here. For further information please refer to `https://lucene.apache.org/core/4_9_0/core/org/apache/lucene/search/similarities/TFIDFSimilarity.html` (last accessed April 2021).

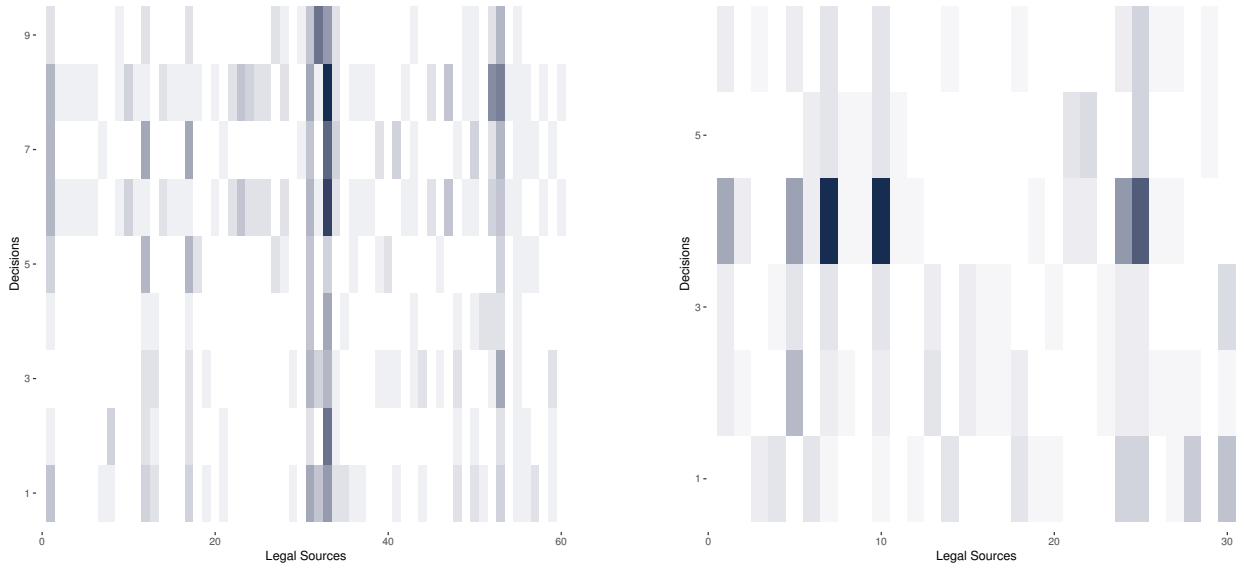

(a) Privacy Infringement —Compensation

(b) Privacy Infringement —Injunction

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

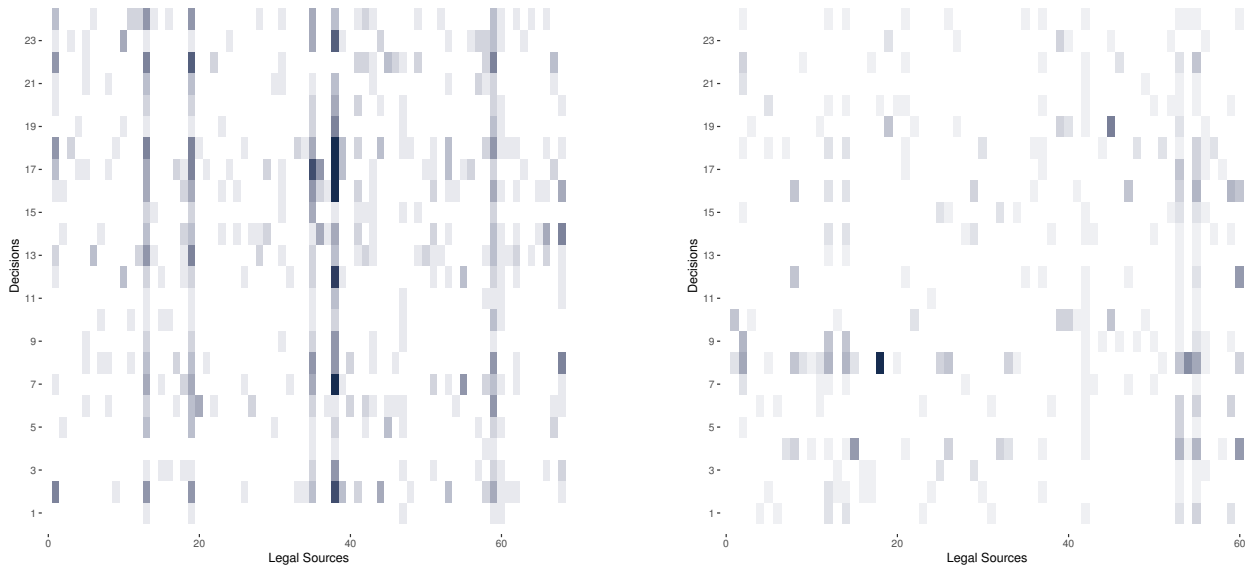

(a) Privacy Infringement —Compensation

(b) Privacy Infringement —Injunction

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 D.1 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-

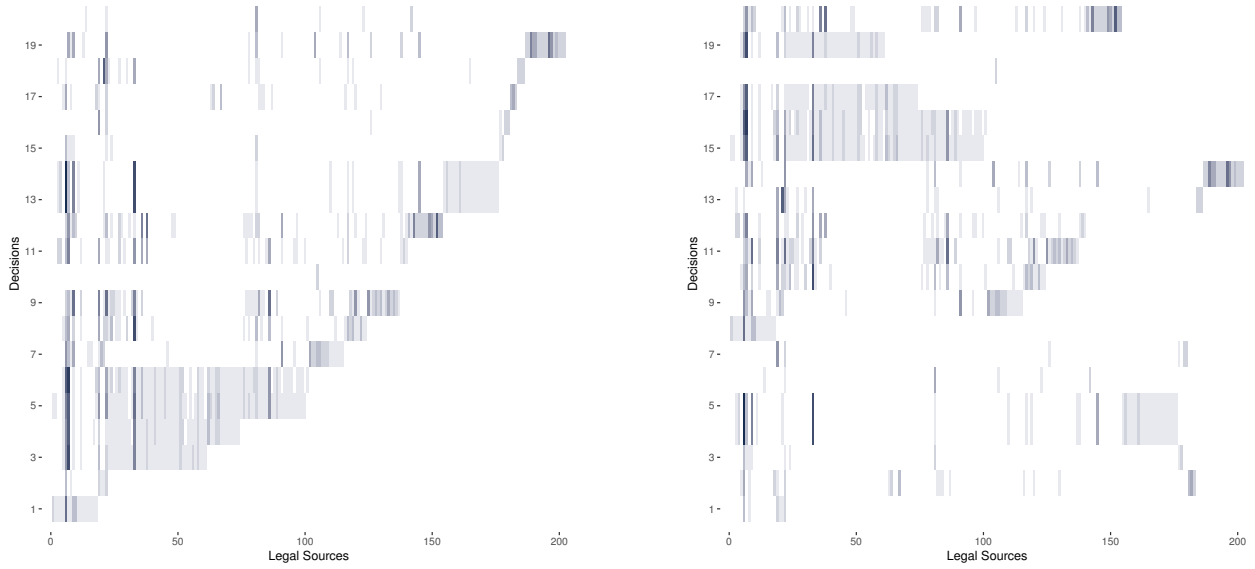

(a) Antitrust —Coding by Occurrence

(b) Antitrust —Coding by Time

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. 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.

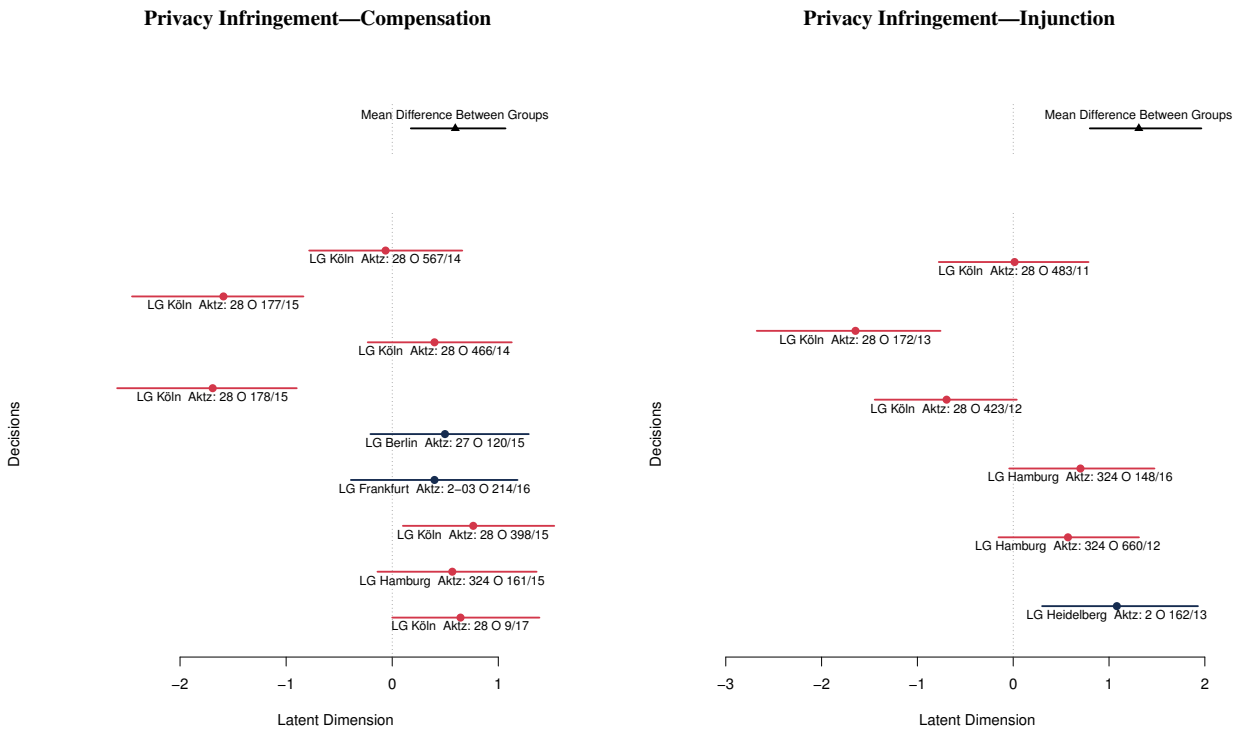

Figure D.1: Estimated Locations of Written Decisions ($\hat{\theta}_i$). Pre-selected Set of Decisions. On the Top: Mean Difference Between the Decisions from Courts in Cologne and Hamburg and All Others. Points Indicate the Median of the Posterior Draws. The Bars Represent the Central 90% Credible Interval.

E Convergence Diagnostics

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

Table E.1: Convergence Diagnostics for Idealpoints. Presslaw Images Exact (Case d_{1A}).

	mean	se_mean	n_eff	Rhat
theta[1]	0.65	0.01	808.69	1.01
theta[2]	0.58	0.01	1,022.59	1.00
theta[3]	0.76	0.02	692.52	1.01
theta[4]	0.40	0.01	1,520.97	1.00
theta[5]	0.53	0.01	1,000.25	1.01
theta[6]	-1.61	0.02	423.78	1.01
theta[7]	0.41	0.01	1,022.99	1.00
theta[8]	-1.52	0.02	440.03	1.01
theta[9]	-0.02	0.01	1,477.29	1.00

Table E.2: Convergence Diagnostics for Idealpoints. Presslaw Images MLT (Case d_{2A}).

	mean	se_mean	n_eff	Rhat
theta[1]	-0.36	0.02	2,312.00	1.00
theta[2]	0.93	0.01	787.16	1.01
theta[3]	0.33	0.01	2,371.99	1.00
theta[4]	0.68	0.02	2,038.15	1.00
theta[5]	-1.11	0.01	1,499.18	1.00
theta[6]	-1.62	0.02	774.71	1.01
theta[7]	1.12	0.01	1,072.90	1.01
theta[8]	-0.54	0.01	1,240.03	1.01
theta[9]	-0.03	0.01	2,943.69	1.00
theta[10]	-0.17	0.01	2,542.65	1.00
theta[11]	0.32	0.01	2,653.62	1.00
theta[12]	1.45	0.01	963.96	1.01
theta[13]	-1.01	0.01	894.82	1.01
theta[14]	-0.29	0.02	683.99	1.01
theta[15]	-0.85	0.01	1,810.35	1.00
theta[16]	0.07	0.01	852.58	1.00
theta[17]	0.47	0.01	915.69	1.01
theta[18]	0.81	0.01	810.71	1.01
theta[19]	0.70	0.01	2,366.71	1.00
theta[20]	0.25	0.01	2,629.76	1.00
theta[21]	-0.29	0.01	2,012.09	1.00
theta[22]	-1.31	0.01	891.44	1.01
theta[23]	1.40	0.01	887.09	1.01
theta[24]	-1.02	0.01	953.60	1.01

Table E.3: Convergence Diagnostics for Idealpoints. Presslaw Online Linking Exact (Case d_{1B}).

	mean	se_mean	n_eff	Rhat
theta[1]	-1.12	0.02	696.96	1.01
theta[2]	-0.59	0.01	1,234.66	1.01
theta[3]	-0.72	0.01	1,210.74	1.01
theta[4]	0.70	0.01	1,064.63	1.00
theta[5]	1.68	0.02	813.58	1.01
theta[6]	-0.02	0.01	1,559.73	1.00

Table E.4: Convergence Diagnostics for Idealpoints. Presslaw Online Linking MLT (Case d_{2B}).

	mean	se_mean	n_eff	Rhat
theta[1]	-0.13	0.01	1,781.65	1.00
theta[2]	-0.26	0.01	1,676.34	1.01
theta[3]	-0.31	0.01	1,662.53	1.00
theta[4]	1.50	0.01	661.85	1.01
theta[5]	0.01	0.01	2,146.56	1.00
theta[6]	-0.30	0.01	1,683.65	1.01
theta[7]	0.01	0.01	1,562.02	1.00
theta[8]	1.23	0.01	665.67	1.01
theta[9]	0.06	0.01	1,364.38	1.01
theta[10]	-1.76	0.01	649.92	1.01
theta[11]	-0.27	0.01	1,995.47	1.00
theta[12]	1.53	0.01	791.62	1.00
theta[13]	0.22	0.01	2,650.93	1.00
theta[14]	-0.17	0.01	1,555.11	1.00
theta[15]	0.49	0.01	1,425.85	1.00
theta[16]	0.85	0.01	943.15	1.01
theta[17]	0.21	0.01	1,595.01	1.01
theta[18]	0.07	0.01	960.50	1.01
theta[19]	-2.08	0.02	612.39	1.01
theta[20]	0.27	0.01	1,397.91	1.01
theta[21]	-0.25	0.01	2,006.41	1.00
theta[22]	0.19	0.01	1,288.17	1.01
theta[23]	-2.06	0.02	788.91	1.01
theta[24]	0.93	0.01	1,225.40	1.00

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

	mean	se_mean	n_eff	Rhat
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 asses 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

[4] We 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 within a case-space the substantive outcomes of decisions show an ordered pattern. In other words, until a certain point courts takes decisions that favor an individual person (in press law) or a cartel victim (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 cartel victim (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 Köln 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 in the case-space on

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

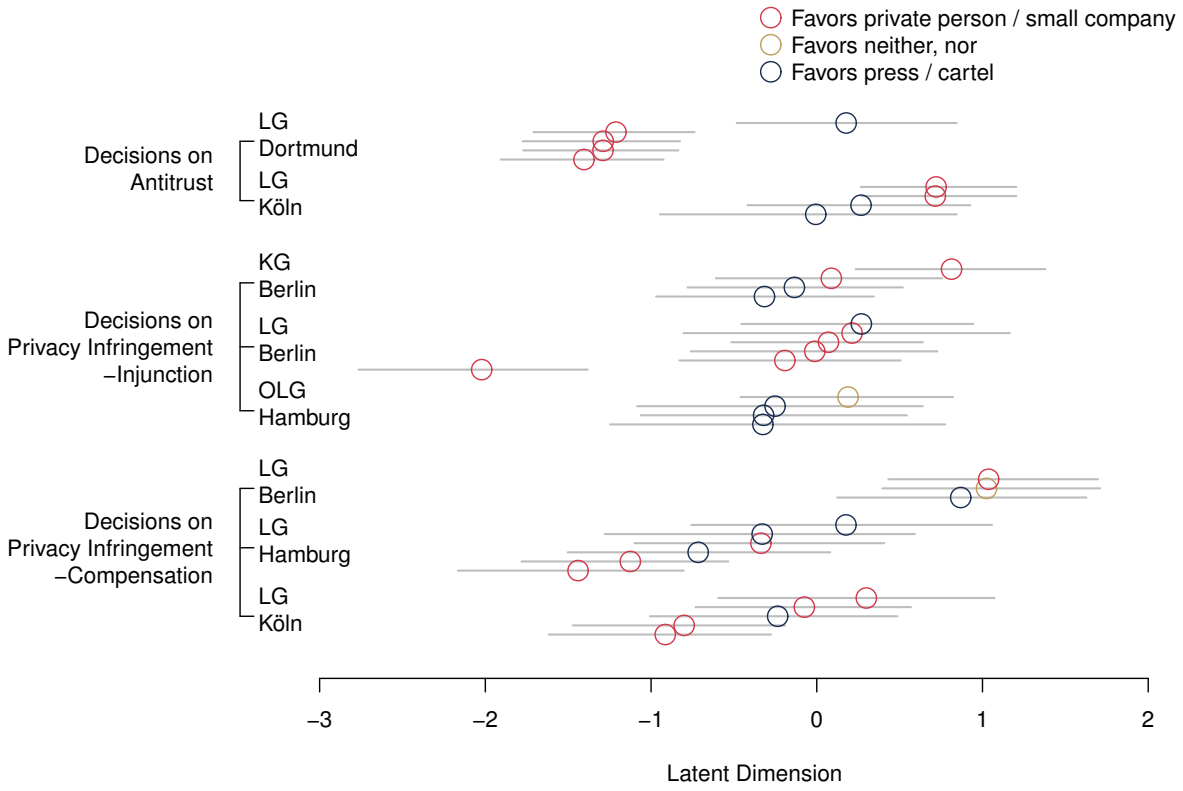

Figure G.1: Estimated Positions of Decisions and their Outcomes. Bars Represent 95% Credible Intervals

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

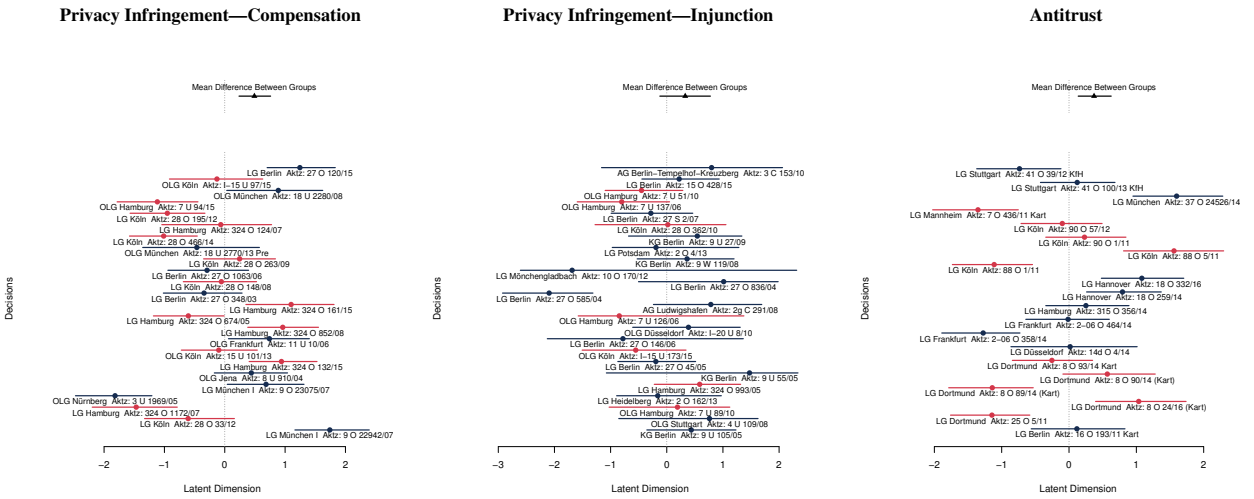

Figure H.1: Positions of Decisions from Randomly Resampled Decision-Source Matrices.

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.

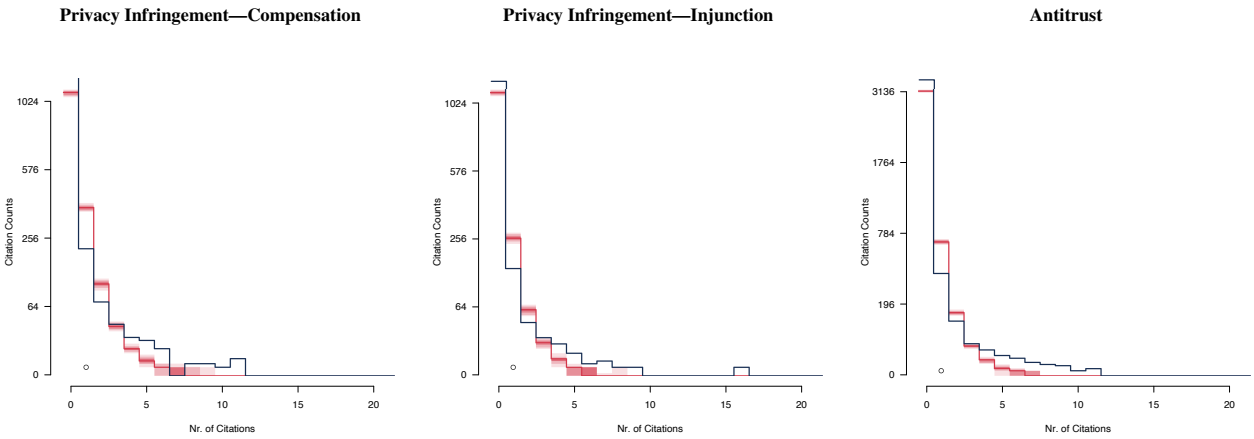

Figure H.2: Posterior Predictive Checks from Randomly Resampled Decision-Source Matrices.