Applied Network Science

RESEARCH   **Open Access**

Quantifying long-term impact of court decisions

CrossMark

Jorge C. Leitão[1,2], Sune Lehmann[1*] (ORCID) and Henrik Palmer Olsen[2]

*Correspondence: sljo@dtu.dk
[1]Technical University of Denmark, Anker Engelundsvej 1, DK-2800 Kgs Lyngby, Denmark
Full list of author information is available at the end of the article

Abstract

In this work, we investigate how court decisions aggregate citations in the European Court of Human Rights. Using the Bass model, we quantify the prevalence of the rich-get-richer phenomenon. We find that the Bass model provides an excellent description of how individual decisions accumulate citations. Our analysis reveals that citations to a large fraction of decisions are, in fact, explained by the rich-get-richer phenomenon. Based on our statistical model, we argue that network properties are insufficient to explain the rich-get-richer effect, suggesting that intrinsic properties of decisions drive a significant part of the observed citation patterns. We conclude by discussing the legal implications of our findings.

Keywords: Law, Citation networks, Bass model, Preferential attachment

Introduction

The advance of computerized systems in the administrative systems of courts enables systematic access to the complete database of their decisions. The use of statistical methodologies to analyze such decisions has brought new insights into the collective institutional behavior of courts. Behaviors that otherwise would not be immediately visible when studying a court's decisions judgment by judgment. In this respect, network science has played an important role in beginning to understand courts' behavior through the analysis of the citation patterns of a number of courts across the globe (Fowler et al. 2006; Fowler and Jeon 2008; Lupu and Voeten 2012; Black and Spriggs 2013).

A citation network of court decisions shares similarities with a network of scientific citations: it is a directed acyclic network whose set of outgoing edges is fixed once the node is created. Furthermore, a citation can be caused by numerous reasons as it can e.g. depend on the semantic content of the two documents, on the semantic content of related items, on the citations by the cited item, etc. Irrespectively of the reason an item is cited, a citation reveals usage of that particular item by the relevant community and the study of this usage reveals important features about the underlying dynamics of the community (Wang et al. 2013; Derlén and Lindholm 2014; Tarissan and Nollez-Goldbach 2015), its structure (Fowler et al. 2006; Fowler and Jeon 2008; Derlén and Lindholm 2017) and, in some cases allows predicting the future usage of individual items (Wang et al. 2013).

A major question in law is how a court uses past decisions to legitimize new decisions (Dworkin 1986). In the case of international courts, where the international law is set by an immutable international treaty, such as the European Convention on Human Rights,

Springer Open

a major question is whether the interpretation of the treaty changes over time, and by how much (Popovic 2008). This process is closely related to how past decisions end up being adopted by the court in future decisions. The European Court of Human Rights has over time interpreted, and continues to interpret, the various rights provisions in a way that can best be described as incrementally progressive. In each case decided, this court tends to either stand firm on its previous decision or to go one step further in its elaboration of what protection for the individual, the convention offers. If it stands firm it will rely on a previous decision as a determinate interpretation of the convention (i.e. not developing the interpretation). If it goes further it will rely on a principle articulated in the earlier decision, and argue that that principle requires that the new case is included under that principle (= developing the interpretation further). In both cases – whether it stands firm or goes one small step further – it will cite it's own previous case law in support of it's decision. In this way the court gradually builds a system of still more elaborate standards and principles, and the citation of earlier decisions plays an important role in this development (Gerards 2018).

In this paper we quantitatively describe the evolution of the usage of court decisions within the European Court of Human Rights (ECHR). Previous research has inquired into legal citations practices in both domestic courts, such as the United States Supreme Court (Fowler et al. 2006) and international courts, such as the Court of Justice of the European Union (CJEU) (Mirshahvalad et al. 2012), the World Trade Organization (WTO) (Pauwelyn 2015), and the International Court of Justice (ICJ) (Alschner and Charlotin 2018). By setting out to quantitatively describe the evolution of the usage of court decisions within the ECHR, our research, therefore, should be seen as a contribution to an emerging body of knowledge on how Courts manage their case-law portfolio via their citation practices.

We propose a model for the usage of the decision by the court based on preferential attachment (Bass model (Bass 1999; 2004)) and show that this model provides an excellent description of the evolution of the number of citations of individual court decisions. We use the Bass model to quantify the extent to which preferential attachment drives the increase in usage of individual decisions – and to predict the future usage of individual decisions by the court. Being able to predict the future usage of a case in terms of how many times it will be cited in the upcoming period, it is possible to identify whether a specific case is on an upward trend or whether its usage has stabilized or is perhaps descending. Combined with knowledge that can be gained from other research (e.g. (Christensen et al. 2016; Šadl and Olsen 2017; Olsen and Küçüksu 2017; Aletras et al. 2016)), about the legal themes and problems dealt with in the case, the prediction about the intensity of future use can be used strategically, for example to guess how the court will reason in the new case.

This paper is structured as follows: in the section "The ECHR and empirical data" we describe the European Court of Human Rights, and the details of the dataset we used. In the section "The use of decisions within the court" we show that the way in which a decision reaches their current number of citations is a heterogeneous dynamical process. In "Dynamical model for the evolution of number of citations" we model each decision's evolution in terms of number of citations as a dynamical process compromising two distinct processes, rich-get-richer and heterogeneous, external factors. We demonstrate the model's usefulness as a description of the vast majority of citation-aggregation in the court's decisions. In "Importance of rich-get-richer effect" we analyze

the prevalence of the rich-get-richer and heterogeneous processes as drivers of the court's decisions and argue that these do not depend on the properties of the network. In the section "Predicting future usage" we use the model to make reliable predictions of the future usage of court decisions, confirming that the model is useful to both describe the past and predict the future. In "Conclusions" we conclude with an outlook of the main results, including a discussion of the legal implications of our findings.

The ECHR and empirical data

The European Court of Human Rights was founded in 1959 and it rules on individual or state applications alleging violations of the civil and political rights set out in the European Convention on Human Rights.

Applications sent to the court can either be admitted or rejected. Some applications are rejected outright as being manifestly ill-founded. Others are admitted or rejected after more careful deliberation that is set out in the court's admissibility decisions. The applications that are admitted to the court will be decided on its merits in a ruling, which will say whether there has been a violation of the Convention. Decisions on the merits are issued as Judgments. It is this final category of decisions (judgments) that we have focused on. Moreover, we only focus on those decisions that have been written in English. While the majority of the cases sent to the court are ruled inadmissible, the cases with an actual judgment are the ones cited by the court in other decisions and therefore these define the court's jurisprudence. For this reason, we restricted this analysis to admitted cases with an English text until the end of 2016. In total, we considered 17 509 cases with 70 861 direct citations between each other. Figure 1 shows the yearly number of these decisions in the database, showing a clear exponential increase of decisions.

The use of decisions within the court

Irrespectively of the content of each decision, past decisions are used by the court as part of its justification for new decisions (see also "Introduction" section above). Operationally, the court does this by citing past decisions. A decision can be cited for different reasons, but, independent of the reason, each citation by definition provides evidence for the use of a previous decision. Therefore, the number of citations is, by definition, a useful quantifier of the usage of a decision within the court (Fowler et al. 2006; Fowler and Jeon 2008; Lupu
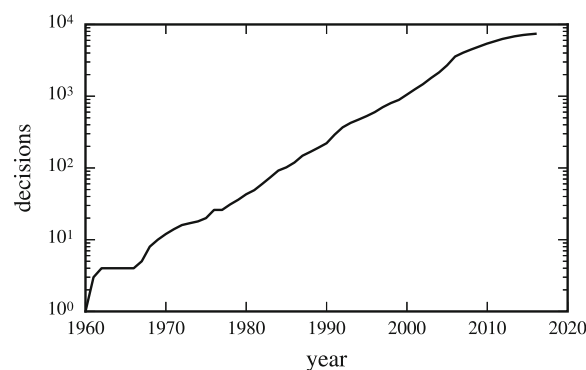

Fig. 1 The exponentially increasing number of decisions by the court from 1960 to present day. The *x*-axis shows time and the *y*-axis shows number of decisions made per year

and Voeten 2012; Black and Spriggs 2013), which we adopt here. The usage of an item is often different from the average usage (Newman 2010); this is also the case here and the resulting broad distribution of citations is also observed in ECHR decisions. Figure 2 (top) shows the ordered (by number of citations) sequence of all decisions of the court until 2016, revealing a broad distribution of the number of citations; some decisions have almost 1000 citations, while a large majority of decisions have zero or very few citations.

There are different possible explanations for the broad distribution of citations. The most well-known is a rich-get-richer mechanism (Simon 1955), also known as preferential attachment (Barabási and Albert 1999). The rich-get-richer phenomenon is present in a variety of natural and cultural contexts, ranging from city population sizes to the number of neighbors in protein-protein interaction networks (Newman 2005; Clauset et al. 2009). In the context of citations between scientific papers, the hypothesis is that the more attention an item has, the more attention it gets.

An important question that arises in the context of the court's citations is whether rich-get-richer is the mechanism that drives the importance of a decision. If yes, why is this the case? What mechanisms contribute for the usage of a decision? To begin to address these questions, we systematically analyze how each highly-cited decision of the ECHR today aggregated their citations over time. As a starting point, we note that the time-evolution of how court decisions are used is extremely heterogeneous (Fig. 3): different decisions take very different paths towards becoming highly used, suggesting that there is no single mechanism which explains the evolution of usage within the court.
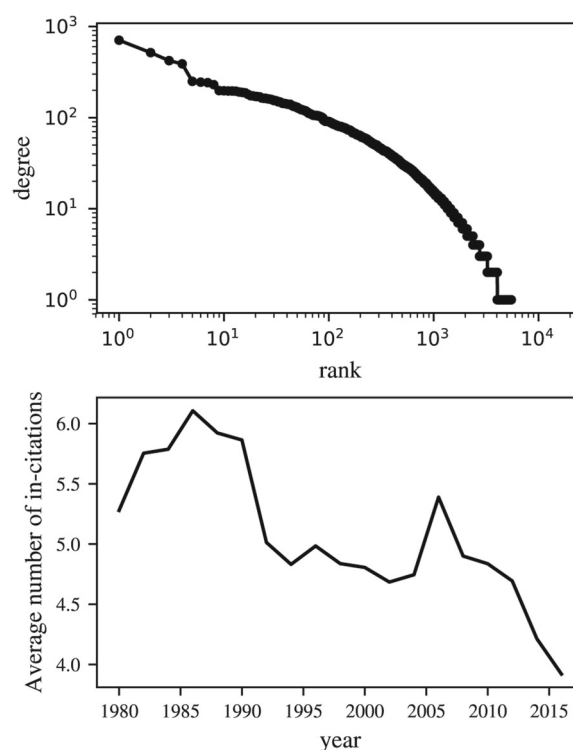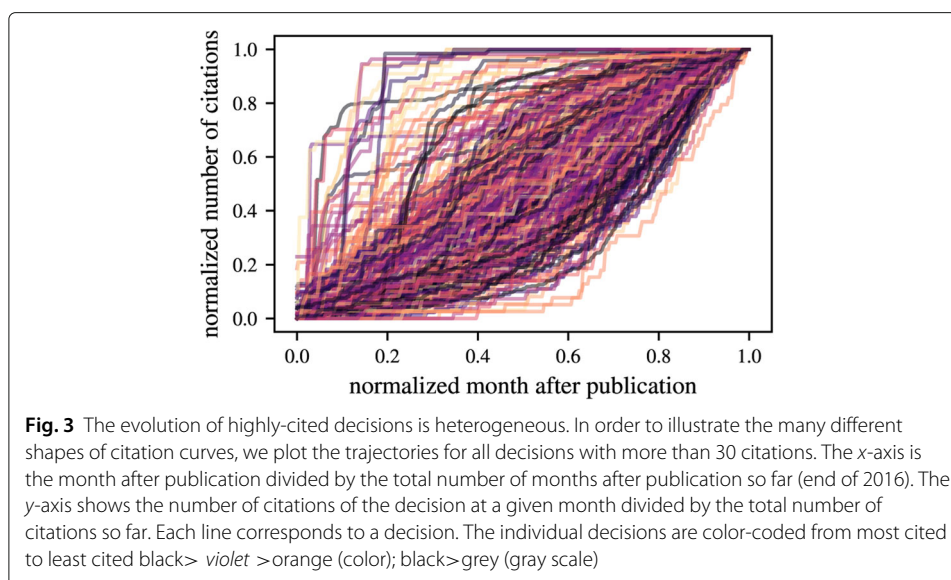

Fig. 2 (top) The ranked number of citations. On the *x*-axis is every decision of the court, where 1 is the most cited decision and the last number the least cited decision. The *y*-axis is the number of citations of the decision. (Bottom) The average number of citations of the decisions as a function of time, showing that a decision will get on average 5 citations, and that this value is roughly independent of time (*y*-axis)

Fig. 3 The evolution of highly-cited decisions is heterogeneous. In order to illustrate the many different shapes of citation curves, we plot the trajectories for all decisions with more than 30 citations. The *x*-axis is the month after publication divided by the total number of months after publication so far (end of 2016). The *y*-axis shows the number of citations of the decision at a given month divided by the total number of citations so far. Each line corresponds to a decision. The individual decisions are color-coded from most cited to least cited black> *violet* >orange (color); black>grey (gray scale)

Dynamical model for the evolution of number of citations

Our first goal is to describe the heterogeneous behaviour presented in Fig. 3 quantitatively. Motivated by the idea that decisions become used/adopted by the court, we consider the Bass model of adoption processes (Bass 1999; 2004). The bass model is one of the simplest models of adoption processes, and has been used to describe different adoption patterns ranging from technology to consumer products, language and citations (Bass 2004). Given its simplicity and ease of interpretation, we consider it to be a relevant model to approach this problem.

The first hypothesis underlying the Bass model is that there is a limit to the number of citations that a decision will aggregate. In the context of law, this corresponds to the notion that, with time, old decisions are less likely to be cited (Tarissan et al. 2016). There are several reasons this might be the case. First of all, the social, economic and political contexts in which a given decision is made will change over time. Even though the legal principle remains the same, at some point the court will find it more appropriate to cite a more recent decision. It will most likely only do so once the principle is so firmly established that it will be seen as uncontroversial. Citing a new case for the same principle may also allow the court to update the principle. By using a slightly changed formulation it may introduce subtle new changes to the principle that will help the court to adapt its case law to changing circumstances (Letsas 2013). Another reason why an old case may not be cited is because the court has refined its case law and thereby the principles it uses to decide cases. What was once one broad principle may have dispersed into two or more more specialized principles, thereby leaving the original principle superfluous or extinct[1]. The speed with which old cases cease to be cited differs according to which specific area of law it belongs. This is because the court receives many cases in some areas and thereby continuously updates and refines the legal principles with which it operates, whereas, in other areas, the court may only get very few cases, which means that old cases may remain important precedent for a much longer time. In some ways this process is similar to how scientific publications (Wang et al. 2013) and technology (Bass 1999) tends to go out of use over time.

The Bass model includes this hypothesis mathematically by stating that the increase in number of citations, $dx(t)/dt$, is proportional to $(c_i - x(t))$ where c_i is a constant representing the maximum number of citations that decision i will obtain. Notice that c_i does not correspond to the maximum number of citations observed in the data, since decisions may still be cited in the future.

The second hypothesis underlying the Bass model is that there are two mechanisms that describe the increase in the number of citations. One mechanism is the *rich-get-richer,* widely known in the community of network science as preferential attachment (Barabási and Albert 1999; Newman 2010). Specifically, this is included in the Bass model in the form of *linear* preferential attachment,

$$\frac{dx(t)}{dt} \propto a_i x(t) (c_i - x(t)) \ ,$$

where a_i is known as the 'fitness parameter' in network science (Newman 2010). This second mechanism generating attention in the Bass model is designed to incorporate *external factors* to the decision (*initial attractiveness* in networks (Newman 2010)), b_i, and describes factors (not proportional to the number of citations) that contribute to the increase of number of citations. These two mechanisms are assumed to be independent. Therefore the full Bass can be written as

$$\frac{dx(t)}{dt} = (b + ax(t)) (c - x(t)), \tag{1}$$

where we have dropped the sub-scripts for simplicity. The solution of this equation (with initial condition $x(0) = 0$) is given by

$$x(t) = \frac{ac \left(e^{t(a+bc)} - 1 \right)}{ae^{t(a+bc)} + bc} \ . \tag{2}$$

To model the behavior observed for the each decision, we fit the parameters a, b and c to the empirical citation curves, using maximum likelihood estimates (see Additional file 1)[2]. In Fig. 4 we show two representative examples of this procedure and respective best estimates.

To confirm the generality of this result, we used two approaches: firstly, we computed the average relative error[3] of the models for every decision: this distribution of average relative errors, Fig. 5), decays exponentially and averages at about 15%. Secondly, we split the data in train (first 80% that we used to fit) and test (last 20% that we used to compute the errors): the distribution of relative errors is similar to the original distribution, see Fig. 9). Overall, these observations demonstrate an excellent representation of a generic decision by the Bass model.

Importance of rich-get-richer effect

We now focus on quantifying the overall importance of each of the two mechanisms described above, namely, rich-get-richer and external factors. We quantify the effect of each mechanism on a given decision by computing the relative contribution, G_{rgr} and G_{ext} to the expected total number of citations c. That is, of all c citations a decision will have, cG_{ext} describes the fraction explained by external factors and cG_{rgr} describes the fraction explained by the rich-get-richer mechanism (according to the model). The fractions G_{rgr}
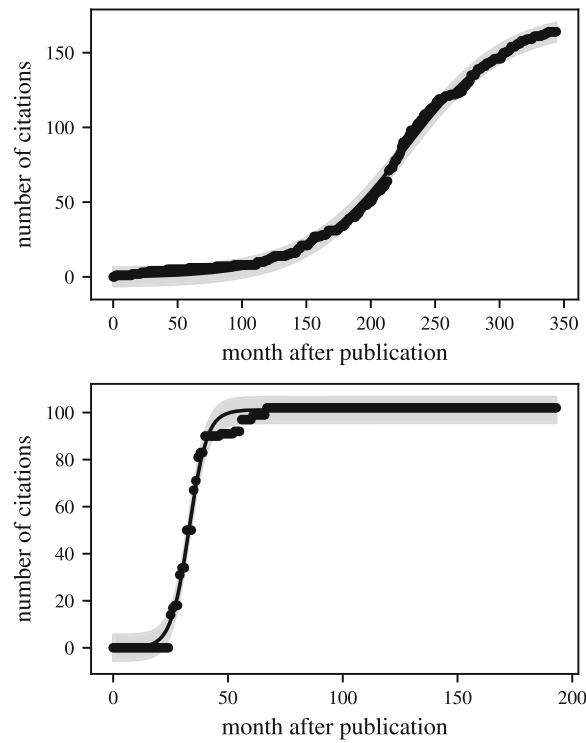
Fig. 4 The best fit of the Bass describes the observed evolution of the number of citations well. The upper panel corresponds to the decision (by merits and just satisfaction) in the case Boyle and Rice V. the United Kingdom (1988); the lower panel corresponds to the decision in the case Edoardo Palumbo V. Italy (2001). The gray region corresponds to the estimated standard deviation (see supp. information)

and G_{ext} can be computed analytically for the Bass model from the parameters a, b and c and are given by Ghanbarnejad et al. (2014)

$$G_{ext} = \frac{a}{bc} \log \left(1 + \frac{bc}{a} \right) \ , \ G_{rgr} = 1 - G_{ext} \ . \tag{3}$$

The relative contributions quantify whether the use of a decision is driven by the rich-get-richer mechanism, external factors, or a combination of the two. A decision with high
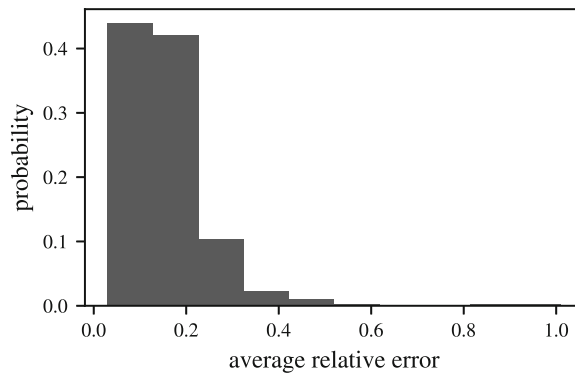

Fig. 5 The empirical distribution of the average relative error ϵ of the model against the data (ensembled over all decisions with more than 30 citations) decays exponentially with increasing error, indicating that the model describes the majority of the cases. The relative error of a decision i is defined as $\epsilon_i \equiv \frac{1}{t^{max}} \sum_{t=0}^{t^{max}} |x_i(t) - x(t)| / x_i(t)$ where $x(t)$ corresponds to Eq. 2 with the best parameters obtained from MLE

G_{rgr} (low G_{ext}) has citations that were obtained gradually throughout time, in a shape that resembles an S curve (see Fig. 4, upper panel for an example). Conversely, a decision with high G_{ext} (low G_{rgr}) characterizes a case where the number of citations were obtained soon after the decision was published, with interest gradually fading (exponential increase after publication, see Fig. 4, lower panel for an example). These contrasting behaviors reflect different ways in which decisions are used by the court.

In Fig. 6 we summarize G_{ext} for all decisions. Analyzing this figure, we note that at around decision rank 50 (approximately 10% of the cases considered) have a value for G_{ext} of 0.5 or larger. The remaining 90% percent of cases are explained primarily by the rich-get-richer mechanism. In fact, for 80% of all decisions, the rich-get-richer mechanism explains more than 70% of all citations. In summary, the majority of citations is well described by a combination of rich-get-richer and external factors, and that the rich-get-richer factors tend to play a larger role in how decisions are cited.

The rich-get-richer mechanism is known to describe well how scientific publications (Wang et al. 2013) and innovations (Bass 1999; 2004) acquire attention. In the context of scientific publications, the reason behind preferential attachment is that the more a scientific publication is known, the more likely it is to be cited, and the more cited it is, the more well known it is.

In legal practice, decisions are decided on the basis of a set of facts specific to the individual case and references should reflect that by specifically referring to earlier cases with a similar fact content. This is the basic content of the classic doctrine *Stare Decisis*[4]. Therefore, it is not immediately obvious that we should expect to see the rich-get-richer effect in legal practice.

We now make a case to why the rich-get-richer should be observed in this court. Our argument is that with increasing case-load and faster paced societal development, legal practices tend to become increasingly standardized and bureaucratized. This leads to a reliance on general principles of law to decide cases. Herman Oliphant noted this development in the US already in 1928 (Oliphant 1928). Oliphant notes a development from *Stare Decisis* to what he calls *Stare Dictis*, on which Judges and Legal Scholars increasingly rely on generalizations from prior cases to form abstract principles of law, which are then more freely applied to new cases.
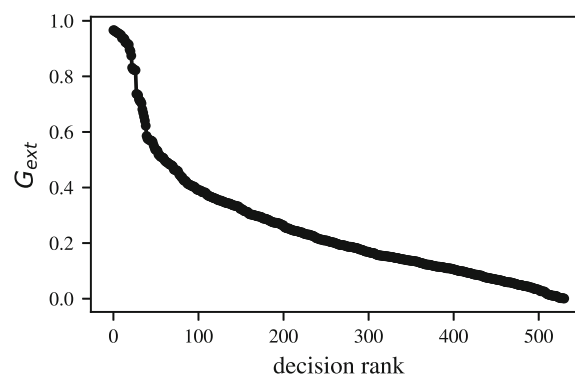

Fig. 6 The rank of all decisions (with more than 30 citations) by G_{ext}. The *x*-axis shows the decision rank, ordered from highest external to lowest external; the *y*-axis shows the value of G_{ext} of the decision

To see how this practice may have emerged in the European Court of Human Rights, we note that the European Court of Human Rights is overburdened with cases, see Additional file 1. To manage its load, the court relies on various procedures that allows it to decide groups of cases that are perceived by the court to derive from the same underlying (structural problem). By definition, this is Stares Dictis: the underlying argument for deciding for the group is no longer the particularities of the individual cases but rather a common abstract principle of law. Since each of these principles are well represented by a limited number of cases, these cases will be competing for their merit in representing the specific principle. This mechanism is well known to lead to a rich-get-richer effect where the more a case is used, the better it represents the principle.

To illustrate how this dynamics is present in the ECHR, we provide examples of two procedures used by the court to mitigate the overburdeness of the court and that results in cases competing for merit. One procedure is the courts' use of 'case law guides'[5] which is aimed at potential litigants and which sets out to explain the main content of the court's case law. In these publications, the court refers to its own case law, another practice which may also enforce the rich-get-richer principle by reinforcing the tendency to use case law highlighted in these publications when deciding future cases. Another procedure is the so-called 'Pilot Judgment Procedure' [6] which allows the court to dismiss a whole group of cases at a time. Those cases that are referenced in the pilot judgments will be considered more important and will therefore feature as leading principles (*Stare Dicta* in Oliphant's terminology) in the Court's case law.

In summary, while one should not immediately expect a rich-get-richer effect in legal practice, the specific practical constraints of the European Court of Human Rights can explain its occurrence. We hypothesise that rich-get-richer is observed in other overburden courts that apply Stare Dictis to mitigate it.

G_{rgr} is an intrinsic property of the decision

Recall that the factors G_{ext} and G_{rgr} are not related to how many citations a decision has (or will have), but how these citations come about. A natural question is why certain decisions are mainly driven by rich-get-richer factors, while others are driven by external effects. We imagine that there is a number of factors that might play a role for a decision to be driven by the rich-get-richer mechanism, such as the number of citations the decision has, the year it was published, how broad the decision is judicially, etc. To understand the role of these different factors, we used a multiple regression model[7] of $y = G_{ext}$ against a broad range of different meta-data and network properties of the decision. Specifically, for each case we considered (for the 2016 citation network)
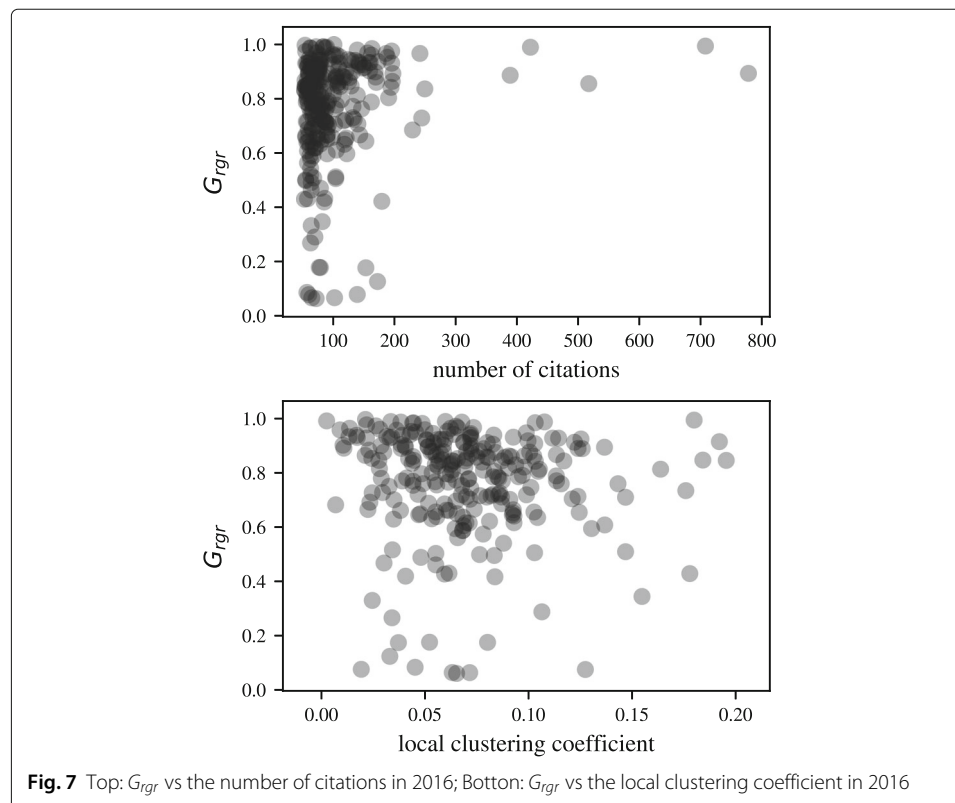
- Number of citations (incoming links).
- The decision's clustering coefficient in the citation network.
- The decision's centrality in the citation network (implemented using Katz, page rank, eigenvector, betweenness, and closeness centrality).
- The year it was published.
- The country the decision is about.
- The outcome of the decision (violation or non-violation of the treaty).
- The judicial breadth of the decision (number of convention articles to which it relates).

- What legal domains the decision relates to (implemented as one-hot encoding of the articles to which the decision relates).

Surprisingly, we found that none of the features listed above have a strong linear correlation with G_{ext} (that is, $R^2 < 0.5$ for each of the features). Further, using the features above, we were unable to formulate a statistical model that explains the observed behaviors. This indicates that G_{ext} (and G_{rgr}) is not a simple consequence of how much the decision was cited, or properties of the network. For example, the fact that the clustering coefficient does not play a role, rules out the hypothesis that high rich-get-richer phenomenon is due to its neighbors being highly connected. Even though no linear correlation exists between G_{rgr} and the measures above, the most cited decisions all have a high G_{rgr} (See top panel of Fig. 7). In other words, a high G_{rgr} is a necessary (but not sufficient) condition for a case to develop into one of decisions most cited by the court. Given our inability to formulate an accurate statistical model based on the features above, our findings suggest that G_{rgr} is an intrinsic property (Wang et al. 2013) of the decision. Such intrinsic properties could be related to legal content and semantic context of a decision rather than the features we have considered here.

Predicting future usage

Knowing which decisions will be used in the future is relevant for all stakeholders of the court. This knowledge can be used to decide when to submit a new case on a specific topic, to identify decisions that are most active in the court, etc. The Bass model describes the evolution of the number of citations of decisions very well; it does so over a large

Fig. 7 Top: G_{rgr} vs the number of citations in 2016; Botton: G_{rgr} vs the local clustering coefficient in 2016
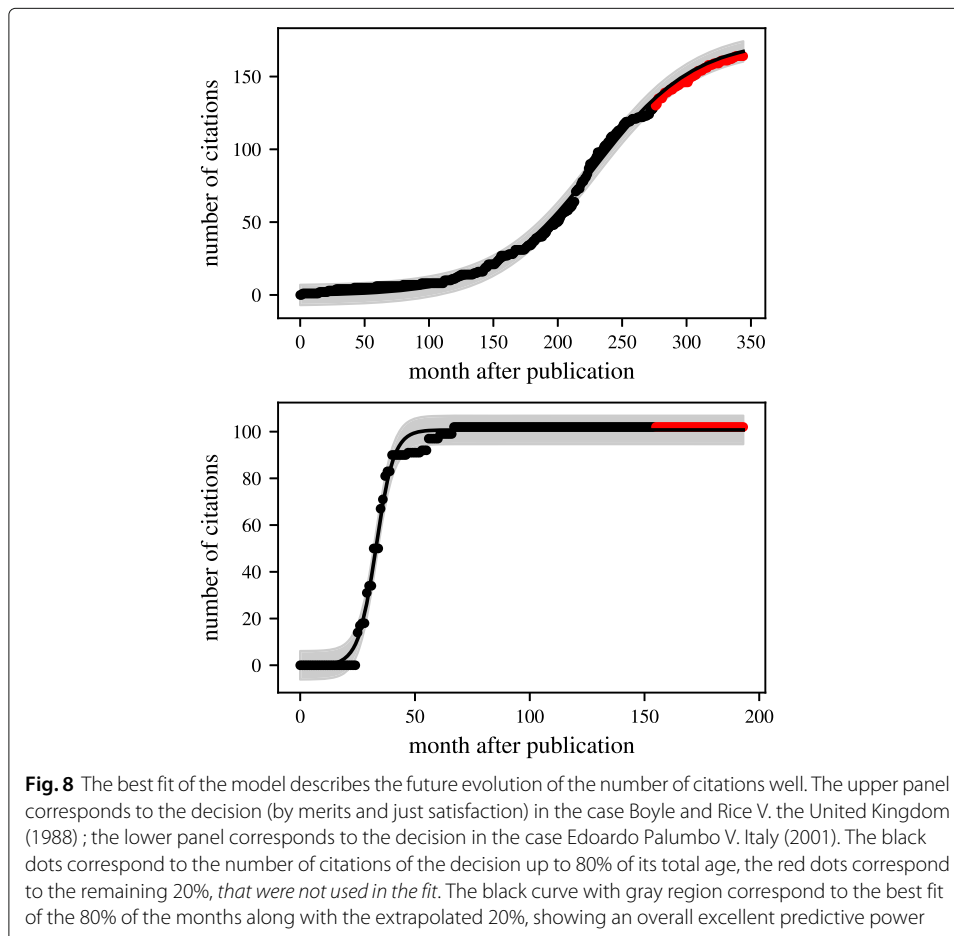
period of time for decisions of different ages. This means that the model can be used to predict decisions' future usage by the court.

To confirm that this is indeed the case, we tested the predictive ability of the model using a simple form of cross-validation. We once again fitted a simple Bass model to each decision. This time, however, we trained (fitted) the model on only 80% of the data available for every decision, while withholding the most recent 20% of data for testing the model. We can now compare the extrapolations based on our training data to the real citations aggregated over the 20% next months. The result of this analysis is shown in Fig. 8 for two examples. To confirm that this holds more generally, we computed the average error over all decisions, which we measured to be less than 10% (see Fig. 9 for the full histogram). This confirms the high predictive power of this model to future usage of court decisions.

Conclusions

In this paper we quantitatively describe the patterns of decisions from one of the most important international courts. Our main findings are that the Bass model provides an excellent description of the time evolution of cases. We confirmed this through two different results: a small RMS error when fitting the model and a overall good predictive power. This model allows us to probe the importance of the rich-get-richer effect in the court's
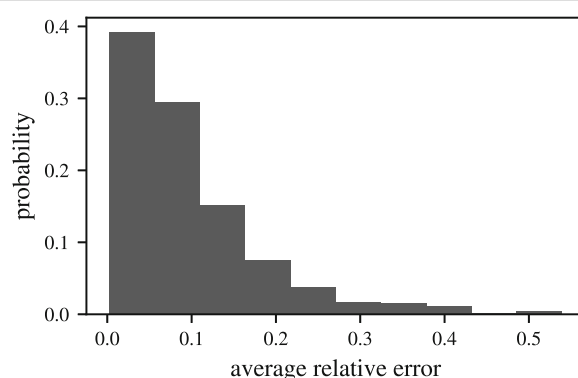

Fig. 8 The best fit of the model describes the future evolution of the number of citations well. The upper panel corresponds to the decision (by merits and just satisfaction) in the case Boyle and Rice V. the United Kingdom (1988) ; the lower panel corresponds to the decision in the case Edoardo Palumbo V. Italy (2001). The black dots correspond to the number of citations of the decision up to 80% of its total age, the red dots correspond to the remaining 20%, *that were not used in the fit*. The black curve with gray region correspond to the best fit of the 80% of the months along with the extrapolated 20%, showing an overall excellent predictive power

Fig. 9 The probability of the average relative error ϵ of the model on the last 20% of hidden months (calculated as an ensemble over all decisions with more than 30 citations) decays exponentially, indicating that the model describes the majority of the cases. The average of this histogram is 10%

citation network, and we find that some citation patterns in the ECHR are characterized by this phenomenon.

The finding that the rich-get-richer effect applies to how decisions are used within the court, is a clear signal that the court has a systematic approach to case citation that goes beyond the traditional *stare decisis* doctrine of fitting the same law to same facts. Previous research has indicated that case citation in very active international courts has become routinized and/or standardized. In Ref. Panagis et al. (2017), for example, the authors show that the CJEU makes use of so called judicial formulae, which are often repeated. Such repeated formulae are used to continuously (re)establish the fundamental principles and concepts of European Union law. Over time, however, the formulae detach from the judgments in which they were first pronounced and acquire a broader relevance. They begin to function as abstract rules and thereby take on a legal significance beyond the fact constellation in which it was first articulated (Panagis et al. 2017). A similar point is made by Ref. McAuliffe (2015), but via a study of how CJEU judgments are crafted[8].

The rich-get-richer observed here effect suggest that the court may have routinized its citation practice, offering the same citations for many decisions in cases that have broadly similar facts. It is an indication, therefore, that the court – in its citation patterns – operates with a broader notion of which citations fit which cases, than one might otherwise expect. The traditional characteristic of judicial decision-making is precisely that it is a very individualized process in which the specific case is scrutinized under the law and decided by carefully comparing it to previous cases that very closely share the same facts. Our finding indicates that the ECHR operates with a broader notion of sameness or similarity, when it cites previous case law as justification for its decisions. This could be seen as a sign that the court operates more freely – and thereby more politically – when it decides cases: There is a larger space between the facts of a case, the justification of the decision, and the citations that is being used to justify that decision. These findings correspond well, with previous legal research which indicated that the court operates a kind of *judicial diplomacy* (Madsen 2010; Olsen 2015).

Finally, to investigate the issue of preferential attachment in further depth, we then attempt to understand how individual cases aggregate citations, using network and contextual (legal domain) features in order to estimate how much the rich-get-richer effect

plays a role for a particular decision. Finally, we find support for explaining the rich-get-richer effect by way of the legal science point-of-view (that we have presented above) in the fact that we were unable to build a useful statistical model to describe the rich-get-richer phenomenon. The features we use do not capture the batching which occurs as part of the 'Pilot Judgment Procedure', nor the 'case law guides' which streamlines the cases themselves. We expect that a promising avenue for future research on this topic could be to explore the connection between these streamlining procedures and citation patterns in the court, a question which has an interplay with the more general challenge of how algorithms shape decision making in public institutions (O'Neil 2016).

Endnotes

[1] As illustrated for example by the court's official case law guide for art. 8, which breaks this short rights provision ("Everyone has the right to respect for his private and family life, his home and his correspondence") up into more than 50 different strings of case law.

[2] We also explored the model proposed in Ref. Wang et al. (2013). This model describes the existing data equally well, but predicts an unrealistically high (10^{10} number of citations for about 30% of all decisions considered), suggesting lack of predictive power in our case.

[3] We used relative error, as opposed to absolute error, because decisions have an skewed number citations, see Fig. 2. A relative metric allows to better interpret the error of the model in respect to the decision's number of citations. A model of a decision whose average relative error is 10% signifies that the model is on average 10% away of each datum

[4] There are numerous public accounts of how this doctrine works in modern law, see for example: https://www.law.cornell.edu/wex/stare_decisis

[5] See: https://www.echr.coe.int/Pages/home.aspx?p=-caselaw/analysis&c=#n1379408479872547532483 7_-pointer.

[6] See: https://www.echr.coe.int/Documents/Pilot_-judgment_-procedure_ENG.pdf.

[7] We also tested a more sophisticated method (a random forest regressor with hyper-parameter optimization and k-fold cross-validation), and found similar results.

[8] Network analysis has also, as mentioned above, been used to investigate case citation behavior in other international courts such as the WTO appellate body and the ICJ. Although network analysis is a useful method for inquiring into case citation behavior in these courts, their case law is not sufficiently rich to fully validate a rich-get-richer phenomenon in these courts. As an illustration of the difference in how much data is produced, Alschner and Charlotin (2018) observes 1 865 citations that the ICJ make to its own case law. In our dataset we observe 70 861 citations that the ECHR makes to its own case law.

Additional file

Additional file 1: Supplementary information. (PDF 36.3 kb)

Authors' contributions
JCL, SL, HPO conceived the paper, JCL prepared and analyzed the data. JCL, SL, HPO wrote the paper. All authors read and approved the final manuscript.

Competing interests
The authors declare that they have no competing interests.

Publisher's Note
Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Author details
[1]Technical University of Denmark, Anker Engelundsvej 1, DK-2800 Kgs Lyngby, Denmark. [2]University of Copenhagen, Karen Blixens Plads 16, DK-2300 Copenhagen, Denmark.

References
Aletras N, Tsarapatsanis D, Preoţiuc-Pietro D, Lampos V (2016) Predicting judicial decisions of the european court of human rights: A natural language processing perspective. PeerJ Comput Sci 2:93
Alschner W, Charlotin D (2018) The growing complexity of the international court of justice's self-citation network. Eur J Int Law 29(1):83–112
Barabási A-L, Albert R (1999) Emergence of scaling in random networks. Science 286(5439):509–512
Bass F (1999) A New Product Growth for Model Consumer Durables. Manag Sci 16(5):1825–1832. https://doi.org/10.1287/mnsc.15.5.215
Bass FM (2004) Comments on "A New Product Growth for Model Consumer Durables The Bass Model". Manag Sci 50(12_supplement):1833–1840. https://doi.org/10.1287/mnsc.1040.0300
Black RC, Spriggs JF (2013) The Citation and Depreciation of U.S. Supreme Court Precedent. J Empir Leg Stud 10(2):325–358. https://doi.org/10.1111/jels.12012
Christensen ML, Olsen HP, Tarissan F (2016) Identification of case content with quantitative network analysis: An example from the ecthr. In: JURIX. IOS Press, Amsterdam. pp 53–62
Clauset A, Shalizi CR, Newman ME (2009) Power-law distributions in empirical data. SIAM Rev 51(4):661–703
Derlén M, Lindholm J (2014) Goodbye van gend en loos, hello bosman? using network analysis to measure the importance of individual cjeu judgments. Eur Law J 20(5):667–687
Derlén, M, Lindholm J (2017) Is it good law? network analysis and the cjeu's internal market jurisprudence. J Int Econ Law 20(2):257–277
Dworkin R (1986) Law's Empire. Harvard University Press, Cambridge
Fowler JH, Jeon S (2008) The authority of Supreme Court precedent. Soc Networks 30(1):16–30. https://doi.org/10.1016/j.socnet.2007.05.001
Fowler JH, Johnson TR, Spriggs JF, Jeon S, Wahlbeck PJ (2006) Network Analysis and the Law: Measuring the Legal Importance of Precedents at the U.S. Supreme Court. Political Anal 15(3):324–346. https://doi.org/10.1093/pan/mpm011
Gerards JH (2018) Dealing with divergence. margin of appreciation and incrementalism in the case-law of the european court of human rights. Hum Rights Law Rev 18(3 1 Sept 2018):495–515. forthcominb Oxford University Press / Oxford Academic. Oxford
Ghanbarnejad F, Gerlach M, Miotto JM, Altmann EG (2014) Extracting information from S-curves of language change. J R Soc Interface 11(101)
Letsas G (2013) The echr as a living instrument: Its meaning and legitimacy. Constituting Eur: Eur Court Hum Rights National, Eur Global Context 2:106
Lupu Y, Voeten E (2012) Precedent in International Courts: A Network Analysis of Case Citations by the European Court of Human Rights. Br J Polit Sci 42(02):413–439. https://doi.org/10.1017/S0007123411000433
McAuliffe K (2015) Behind the scenes at the court of justice: A story of process and people. In: Davies B, Nicola F (eds). EU Law Stories. Cambridge University Press, Cambridge. pp 141–172
Madsen MR (2010) Legal diplomacy: law, politics and the genesis of postwar european human rights. In: Human Rights in the Twentieth Century. Cambridge University Press, Cambridge. pp 62–81
Mirshahvalad A, Lindholm J, Derlén M, Rosvall M (2012) Significant Communities in Large Sparse Networks. PLoS ONE 7(3):33721. https://doi.org/10.1371/journal.pone.0033721
Newman ME (2005) Power laws, pareto distributions and zipf's law. Contemp Phys 46(5):323–351
Newman M (2010) Networks: An Introduction. Oxford University Press
Oliphant H (1928) A return to stare decisis. ABAJ 14:71
Olsen HP (2015) International courts and the doctrinal channels of legal diplomacy. Transl Leg Theory 6(3-4):661–680

Olsen HP, Küçüksu A (2017) Finding hidden patterns in ecthr's case law: On how citation network analysis can improve our knowledge of ecthr's article 14 practice. Int J Discrimination Law 17(1):4–22

O'Neil C (2016) Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Broadway Books, Danvers

Panagis Y, Sadl U, Tarissan F (2017) Giving every case its (legal) due the contribution of citation networks and text similarity techniques to legal studies of european union law. In: 30th International Conference on Legal Knowledge and Information Systems (JURIX'17), vol 302. IOS Press, Amsterdam. pp 59–68

Pauwelyn J (2015) Minority rules: Precedent and participation before the wto appellate body. In: Jemielniak J, Nielsen L, Olsen HP (eds). Judicial Authority in International Economic Law. Cambridge University Press, Cambridge. pp 141–172

Popovic D (2008) Prevailing of judicial activism over self-restraint in the jurisprudence of the european court of human rights. Creighton L Rev 42:361

Šadl U, Olsen HP (2017) Can quantitative methods complement doctrinal legal studies? using citation network and corpus linguistic analysis to understand international courts. Leiden J Int Law 30(2):327–349

Simon HA (1955) On a class of skew distribution functions. Biometrika 42(3/4):425–440. Kettering

Tarissan F, Nollez-Goldbach R (2015) Temporal properties of legal decision networks: a case study from the international criminal court. In: 28th International Conference on Legal Knowledge and Information Systems (JURIX'2015). IOS Press, Amsterdam

Tarissan F, Panagis Y, Šadl U (2016) Selecting the cases that defined europe: complementary metrics for a network analysis. In: Advances in Social Networks Analysis and Mining (ASONAM), 2016 IEEE/ACM International Conference On. IEEE, Piscataway. pp 661–668

Wang D, Song C, Barabási A-L (2013) Quantifying Long-Term Scientific Impact. Science 342(6154):127–132. 10.1126/science.1237825