This article was downloaded by: [149.157.61.223] On: 16 June 2025, At: 08:35 Publisher: Institute for Operations Research and the Management Sciences (INFORMS) INFORMS is located in Maryland, USA



Information Systems Research

Publication details, including instructions for authors and subscription information: $\underline{http://pubsonline.informs.org}$

User-Generated Content Shapes Judicial Reasoning: Evidence from a Randomized Control Trial on Wikipedia

Neil C. Thompson; , Xueyun Luo; , Brian McKenzie; , Edana Richardson, Brian Flanagan

To cite this article:

Neil C. Thompson: , Xueyun Luo; , Brian McKenzie; , Edana Richardson, Brian Flanagan (2024) User-Generated Content Shapes Judicial Reasoning: Evidence from a Randomized Control Trial on Wikipedia. Information Systems Research 35(4):1948-1964. <u>https://doi.org/10.1287/isre.2023.0034</u>

This work is licensed under a Creative Commons Attribution- NonCommercial 4.0 International License. You are free to download this work and share with others for any purpose, except commercially, and you must attribute this work as "*Information Systems Research*. Copyright © 2024 The Author(s). <u>https://doi.org/10.1287/isre.2023.0034</u>, used under a Creative Commons Attribution License: <u>https://creativecommons.org/licenses/by-nc/4.0/</u>."

Copyright © 2024 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

User-Generated Content Shapes Judicial Reasoning: Evidence from a Randomized Control Trial on Wikipedia

Neil C. Thompson,^{a,*} Xueyun Luo,^b Brian McKenzie,^c Edana Richardson,^d Brian Flanagan^d

^a MIT Computer Science & Artificial Intelligence Laboratory, MIT Initiative on the Digital Economy, Cambridge, Massachusetts 02139; ^b SC Johnson College of Business, Cornell University, Ithaca, New York 14853; ^c Critical Skills Programme, Maynooth University, Maynooth, W23 F2H6 County Kildare, Ireland; ^d School of Law and Criminology, Maynooth University, Maynooth, W23 F2H6 County Kildare, Ireland; ^{*}Corresponding author

Contact: neil_t@mit.edu, (b https://orcid.org/0000-0001-9888-573X (NCT); x1777@cornell.edu, (b https://orcid.org/0000-0003-2853-1033 (XL); brian.mckenzie@mu.ie (BM); edana.richardson@mu.ie, (b https://orcid.org/0000-0003-0569-5481 (ER); brian.flanagan@mu.ie, (b https://orcid.org/0000-0003-0569-5481 (ER); b https://orcid.org/0000-0003-0569-548

Received: January 19, 2023 Revised: September 10, 2023; December 6, 2023 Accepted: December 24, 2023 Published Online in Articles in Advance: March 19, 2024

https://doi.org/10.1287/isre.2023.0034

Copyright: © 2024 The Author(s)

Abstract. Legal professionals have access to many different sources of knowledge, including user-generated Wikipedia articles that summarize previous judicial decisions (i.e., precedents). Although these Wikipedia articles are easily accessible, they have unknown provenance and reliability, and therefore using them in professional settings is problematic. Nevertheless, Wikipedia articles influence legal judgments, as we show using a first-of-its-kind randomized control trial on judicial decision making. We find that the presence of a Wikipedia article about Irish Supreme Court decisions makes it meaningfully more likely that the corresponding case will be cited as a precedent by judges in subsequent decisions. The language used in the Wikipedia article also influences the language used in judgments. These effects are only present for citations by the High Court and not for the higher levels of the judiciary (Court of Appeal and Supreme Court). The High Court faces larger caseloads, so this may indicate that settings with greater time pressures encourage greater reliance on Wikipedia. Our results add to the growing recognition that Wikipedia and other frequently accessed sources of usergenerated content have profound effects on important social outcomes and that these effects extend farther than previously seen-into high-stakes settings where norms are supposed to restrict their use.

History: Rajiv Kohli, Senior Editor; Gordon Burtch, Associate Editor.

Open Access Statement: This work is licensed under a Creative Commons Attribution- NonCommercial 4.0 International License. You are free to download this work and share with others for any purpose, except commercially, and you must attribute this work as "Information Systems Research. Copyright © 2024 The Author(s). https://doi.org/10.1287/isre.2023.0034, used under a Creative Commons Attribution License: https://creativecommons.org/licenses/by-nc/4.0/."

Funding: The authors thank the Massachusetts Institute of Technology and Maynooth University for research funding.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/isre.2023.0034.

Keywords: Wikipedia • user-generated content • knowledge diffusion • law and economics • randomized control trial

1. Introduction

The flourishing of user-generated content (UGC) in recent decades has sparked interest in who is creating this content and what influence it is having (Forman and van Zeebroeck 2019). In another age, the influence of content created by a random person might have been confined to their real-world friends and colleagues. However, since the emergence of digital platforms in the early 2000s, the creation and reach of UGC has expanded dramatically (Fader and Winer 2012). One of the most important UGC platforms is Wikipedia, which has more than 6.5 million articles in English (Wikipedia 2022) and gets 18 billion page views per month (Pew Research Center 2016). Even these figures undervalue Wikipedia's

influence, because search engines also surface Wikipedia content in "knowledge panels" (to the right of search results) without users ever visiting the site.

Wikipedia has already been shown to influence many types of informal decisions. We consider the opposite: Whether Wikipedia influences judicial reasoning, an area of knowledge where decision making is supposed to be a highly structured combination of expert judgment and precedent. This contrasts with existing research that quantifies Wikipedia's causal effects on real-world outcomes that lack formal processes, such as the destination choices of tourists (Hinnosaar et al. 2023), what scientists publish (Thompson and Hanley 2018), and financial investment decisions (Xu and Zhang 2013). One might expect formal modes of decision making, like law, to be less susceptible to the effects of UGC. In fact, we show that UGC matters there as well.

We establish a causal connection between Wikipedia and judicial decisions by conducting a randomized control trial, in which summaries of 154 Irish Supreme Court cases were written up as Wikipedia articles and then randomized into two groups: 77 treatment case summaries were uploaded to Wikipedia and 77 control case summaries were not made public. We then looked to see whether the randomly chosen Supreme Court cases where summaries were added to Wikipedia played a larger role in judicial decision making than did those cases without Wikipedia summaries. They did. Although the evidence we present comes from the Irish legal system, we argue that these findings are likely to apply to other jurisdictions, including the United States.

Our results establish several facts about how Wikipedia influences legal judgments. First, we observe that adding a Wikipedia article about a Supreme Court case increases the frequency with which that case is cited in future decisions. This effect is concentrated in courts that have "first instance" jurisdiction, where legal adjudication begins and where case volumes are highest. In contrast, we find that Wikipedia has no statistically significant influence on decisions made by the Court of Appeal or by the Supreme Court (higher courts that handle appeals). Second, we observe that the contextualization provided in the text of the Wikipedia articles is echoed in the text of judicial opinions. Thus, we find that Wikipedia is influencing both the legal authorities to which the judgment appeals and the form of the legal argument it presents. Last, we explore the mechanisms through which these effects occur, where evidence supports the presence of multiple channels of influence, including via judges (and their staffs), lawyers, or other legal professionals.

Overall, our paper makes three major contributions to the existing literature, which we outline here and expand on in the discussion section. First, we provide causal evidence of the role of UGC in legal reasoning. There has been evidence of UGC affecting nonprofessional decisions, such as the destination decisions of tourists (Hinnosaar et al. 2023) and financial investment decisions (Xu and Zhang 2013). A few studies have discussed the influences of UGC on professional practice, such as the knowledge diffusion by online discussion forums among software development community (Huang et al. 2022), impact on journalists (Sen et al. 2023), and what scientists publish (Thompson and Hanley 2018). Our study, unlike prior work, shows the impact of UGC extends to the context of law, a setting that is highly regulated, professionalized, and is itself

of substantial public interest because of the foundational role of law in society.

Second, our study sheds new light on the implied conflict between UGC's availability and authoritativeness, helping to make clear the limits of UGC's influence. On the one hand, it has long been shown that easily available sources of information influence followon knowledge production and diffusion (Phillips et al. 1991, Ahn et al. 2016), and it might be easy to conclude from these findings that Wikipedia should have influence everywhere regardless of professionalism. On the other hand, professional norms attenuate UGC's value in the contexts where strict procedure is supposed to be followed (Hildebrandt 2016). Okoli et al. (2014) argue that Wikipedia continues to operate primarily as a source of general information for the public rather than as a resource for or site of engagement by professionals. Our results show that the truth lies in the middle. Normative prohibitions do seem to keep Wikipedia from influencing the most-important, well-resourced parts of law, but these prohibitions are insufficient when time and resources pressures would incentivize the use of Wikipedia. Hence our research helps delineate the boundaries of UGC's influence on professional decision making

Finally, by showing that Wikipedia can exert influence in as important and formal an area as law, our paper reinforces the case for accurate, reliable UGC (Kane and Ransbotham 2016, Greenstein and Zhu 2018), especially in domains with potentially far-reaching societal consequences. Given that there is no clear way to prevent individuals from making use of UGC professionally or nonprofessionally, our findings also contribute to the ongoing discussion of how to build public repositories of knowledge, such as Wikipedia, into more reliable storehouses (Chhabra and Iyengar 2020).

2. Legal Context

Although law is a highly professionalized domain, Wikipedia has long been used openly by judges to answer "questions of fact," that is, questions about the nature of the situation to which the law is to be applied.¹ Citations to Wikipedia in U.S. judicial opinions first appeared in 2004 (Peoples 2009). By 2010, there were at least 117 U.S. state and 326 federal cases that cited Wikipedia (Gerken 2010). By October 2022, a Westlaw search indicated that Wikipedia had been referenced in 545 U.S. state and 1,834 federal cases. In this study, we consider whether Wikipedia is also being used, implicitly, to answer "questions of law," that is, questions about the content of the law to be applied. If the information on Wikipedia is of dubious or unknown credibility, then such use might be at odds with litigants' entitlement to an expert resolution of the legal aspect of their dispute: "The rule of law ... requires that citizens receive predictable and nonarbitrary treatment when they seek relief in the judicial system" (Hitt 2019, p. 82). By contrast, per their own description: "Wikipedia allows anyone with an Internet connection to alter its content. Please be advised that nothing found here has necessarily been reviewed by people with the expertise required to provide you with complete, accurate, or reliable information. The content of any given article may recently have been changed, vandalized, or altered by someone whose opinion does not correspond with the state of knowledge in the relevant fields If you need specific advice (for example, medical, legal, financial, or risk management), please seek a professional who is licensed or knowledgeable in that area" (Wikipedia 2023).

The setting for our experiment is the Irish legal system, which offers two advantages. First, at the time of our experiment there were fewer than 10 Wikipedia articles on Irish legal cases, which allowed us to do a larger intervention: We could write more articles and could choose particularly consequential cases to write about. Had we instead chosen a jurisdiction like the United States, where many Supreme Court cases already have Wikipedia articles, our intervention would have had to focus on less influential cases and would have required a much longer trial to detect any effects. Second, the Irish legal system is appealing because of its similarity to many other legal systems around the world. Ireland, like England, the United States, much of Canada, and many other countries, has a common-law legal system in which lower courts are bound to follow higher court decisions and where judges routinely cite earlier decisions as precedents that support their conclusion (Berman 2000, Schiavoni 2002). Judicial behaviors in the Irish legal system, such as those documented by our experiment, could be a feature of common-law jurisdictions more broadly. In what follows, we focus on the similarities between the Irish and American legal systems.

In a common-law system, legal rules are articulated, shaped, and developed by judges in the context of specific court cases. Although the legislative arm of the state enacts laws in the form of legislation, the judiciary applies those laws in individual situations and thus helps to determine their practical scope and limitations. In performing this function, the Irish courts share a similar hierarchical structure to that found in the United States. Sitting at the apex of this hierarchy, the Irish Supreme Court is the court of final appeal. As a constitutional court, the Irish Supreme Court is the final arbiter on the interpretation of the Irish Constitution and has jurisdiction to review any law for conformity with the Constitution. Ireland's Supreme Court is also relatively active: In 2019, it delivered 130 written judgments on the merits (of the 272 applications for leave to appeal from the High Court and Court of Appeal in

that year). Sitting below the Irish Supreme Court, the Irish Court of Appeal was established in 2014 and is an appellate court that has both civil and criminal divisions. The Court of Appeal will hear appeals from the High Court, which is both a court of first instance in civil and criminal matters that are beyond the jurisdiction of the lower courts, and an appellate court in civil matters and in certain family law matters.²

As in the United States, the doctrine of precedent (*stare decisis*) is central to the Irish legal system and dictates that a court is bound by earlier court decisions on analogous legal issues (Lobingier 1946). As such, the decisions of the Supreme Court are binding on all inferior Irish courts, including the Court of Appeal and the High Court. Court judgments will therefore include citations to previously decided cases to explain how the decision in the relevant case was reached based on precedent.³

Once a decision is handed down in an Irish court case, open access to written judgments is generally available through the website of the Irish Courts Service, which publishes written judgments issued by the High Court, Court of Appeal, and Supreme Court.⁴ This website contains virtually all written Supreme Court judgments issued since 2001 and Court of Appeal and High Court judgments issued since 2014 and 2005, respectively. (All written judgments of the Irish courts are considered published and of precedential value.) Once published on the Irish Courts Service website, written judgments are collated, organized, and indexed by subscription-based commercial legal databases.

In the course of litigation, each litigant, through their legal team, submits a legal brief to the court designed to show that relevant statutory provisions, precedents, and other legal materials support their preferred view of the applicable law. Together, the parties' lawyers also submit to the court a jointly compiled "book of authorities" that consists of copies, and often summaries, of some of the previous judgments mentioned in both competing legal briefs.⁵ When looking for a summary of a precedent mentioned in a litigant's legal brief, a judge (or their clerk) might search for it in a specialist legal resource such as the submitted book of authorities or in a professional database such as vLex. A potentially more efficient alternative would be to use an Internet search.

On entering a general Internet search (e.g., using Google, Bing, or DuckDuckGo) of a legal judgment, summary information in the form of a "knowledge panel" is often prominently displayed to the right of the search results. The knowledge panel routinely includes information from a corresponding Wikipedia article and a link to that entry. The Wikipedia article will often also show up near the top of the search results. In this way, Wikipedia offers greater expedience than any source of professionally produced judgment summaries because it immediately provides conspicuous summary information (in the form of the knowledge panel's text, together with follow-on links) via the very same search function with which the judge/assistant conducts all their Internet searches. As such, Wikipedia connects with judges' generic process of information retrieval. Existing research indicates that judges with heavier workloads tend to write shorter judgments and to produce more verbal (ex tempore) judgments (Engel and Weinshall 2020). It stands to reason that, under work pressure, Wikipedia's expedience might also make its case summaries an attractive resource for legal professionals.

3. Experiment

Over the period 2019-2020, we conducted a randomized field experiment.⁶ We created 154 new Wikipedia articles, authored by law school faculty and students, that summarized Irish Supreme Court cases. Half of these cases (77) were randomly chosen to be the treatment group, and the summaries were uploaded to Wikipedia for use by legal practitioners and the general public. These Wikipedia articles were added in three waves: seven articles were added in April 2019, four articles were added in October 2019, and the remaining 66 articles were added in May 2020. The other 77 Supreme Court cases represented the control group, and those summaries were not uploaded to Wikipedia and instead served as the counterfactual for what would have happened to the treatment group had it not been uploaded to Wikipedia. After uploading the summaries for the treatment cases to Wikipedia, we observed as subsequent legal cases cited these treatment cases and whether that differed from the control cases. We also monitored the uploaded Wikipedia articles for edits,

but there were no substantive content changes, only minor copyediting and formatting ones.

3.1. Experiment Design

For the experiment, we made choices both for which Irish Supreme Court cases to include and how to write up their Wikipedia summaries. Our goals were to maximize our statistical power (and therefore our chances of detecting an effect if it really existed) and to minimize the presence of confounders that might mislead the analysis. A summary of this process is shown in Figure 1.

We chose 7 areas of law of the 26 present in JustisOne⁷: administrative and constitutional law; asylum, immigration and nationality; banking and finance; crime and sentencing; family law; practice and procedure; and tort. Based on our analysis of a sample of historical data, these categories offered two benefits. First, they had more Supreme Court decisions, which facilitated finding pairs of decisions of similar vintage to stratify together. Second, they had more citations per decision, making it more likely that we would be able to detect effects.

Within these seven areas of law, we chose decisions that would minimize the nonexperimental variation between the treatment and control groups, thereby maximizing our ability to detect an experimental effect. We did this by finding nearest neighbor cases that were highly alike using the *quickblock* library from R. We deemed two cases to be nearest neighbors if they were (1) in the same area of law and decided in the same year and (2) maximally similar in their number of positive citations (in support of judge's conclusion), neutral citations, negative citations, RTE Ireland (TV channel) references, *Irish Times* (newspaper) references, other media references, and publication year (because this is

Figure 1. (Color online) Wikipedia Article Creation Process

1 Population	2 Sample	3 Matching	4 Blocks	5 Randomization	6 Intervention
Irish Supreme Court Cases	Areas of law:Administrative	• Area of law	Written by one author		Treated cases (77): Wikipedia
	and constitutional law asylum	Decision year			Control cases (77): Not made public
	 Immigration and nationality 	Nearest Neighbor:			
	Banking and finance Crime and sentencing	# of positive/ neutral/negative citations			
	Family law	newspaper citations			
	Practice and procedure	other media references			
	Tort	 publication year 			

not always the same as the judgment year). This matching was done by minimizing the Mahalanobis distance after each similarity variable was studentized and given equal weight (except for the publication date, which was weighted at five times). This produced case pairs that were about different topics but quite similar in terms of their vintage, type of law, and citation behavior. That is, before treatment, they behaved comparably in terms of citations (our key outcome), despite being on distinct topics. (To illustrate, we provide details of two of our blocks in the online appendix.) That is, we are accomplishing ex ante by design what propensity score matching attempts to do ex post.

Having grouped our cases in similar pairs, we had students write them up into Wikipedia articles. In 2019 and 2020, undergraduate and graduate law students at Maynooth University each selected a pair of articles from one of the selected categories of law. They were provided with electronic resources to guide their article writing, an induction session, and ongoing editing support from law faculty at Maynooth University, as well as exemplar Wikipedia case articles prepared by law faculty.⁸ For our study, it was important that both groups of articles were written up because it meant that a linguistic analysis could be conducted to determine whether Wikipedia influences judicial opinions' language.

We tracked the outcomes for the experiment (citations, language of judicial decisions) using data from vLex Justis and the court.ie website. Because of their dynamic interfaces and changing ownership, significant manual downloading and reconciliation were necessary to ensure a clean, correct data set. The summary statistics for our chosen articles are shown in Table 1.

From each pair of drafted articles, we randomly assigned one to treatment and one to control. Because of our experimental design, with nearest neighbor stratification to minimize variation on observed characteristics and randomization to minimize variation on unobserved characteristics, we expected our treatment and control groups to be highly similar.

3.2. Randomization Check and Covariate Balance

To test the success of our randomization in producing comparable treatment and control groups, we check for covariate balance pretreatment. Figure 2(a) shows the distributions in the number of citations made by judges to each set of cases per month in the pretreatment period. It reveals that the mean citations per month are similar between the treated group and control group. In addition, a full count of the pretreatment citations to the treatment and control groups also shows balance, Figure 2(b). Both figures provide convincing evidence that the randomization produced treatment and control groups with similar citation behavior.

In addition, we check whether the covariates we used in our randomization process are indeed similar between the treated and control cases. As shown in Table 2, they are indeed similar. Finally, as reported later, we check for parallel trends in the pretreatment citation patterns of treatment and control cases and find them to be very similar.

4. Empirical Analysis and Results

In our analysis, we consider the month of the intervention as t = 0, with periods beforehand ("pretreatment period") or after ("posttreatment") labeled negatively or positively based on the months since that time. The end of our sample period is on December 31, 2020, which corresponds to a different *t* for each wave of articles. In total, we observed more than 7,200 citations to our sample cases.

The analysis that follows is in line with our preregistered plan,⁹ with the following caveats that arose because of data constraints. The pandemic precipitated a reduction in the number of cases decided by the Irish judiciary and hence in the number of new citations to previous judgments (Courts Service Annual Report

 Table 1. Summary Statistics of Irish Supreme Court Cases in Our Sample

Type of Law % of Sample	Tort 5.2	Criminal 16.9	Family 10.4	Finance 5.2	Immigration 14.3	Procedural 32.5	Constitutional 15.6
Variables		Ν	Mean	Standard deviation	Median	Minimum	Maximum
Decision year		154	2009	5.4	2009	2000	2017
Age		154	10.7	5.4	11	3	20
No. of positive citations		154	3.0	6.3	1	0	56
No. of neutral citations		154	22	26	13	0	199
No. of negative citations		154	0.42	0.86	0	0	5
No. of Irish Times citations		154	0.43	0.91	0	0	9
No. of RTE citations		154	0.08	0.28	0	0	1





Notes. (a) Average number of citations over time in the pretreatment period. (b) Distribution of total number of citations in the pretreatment period.

2020). In light of this development, it became inappropriate to use log(citations) as the functional form of the dependent variable because log(0) is undefined and the typical log(citations + 1) approximation is highly inaccurate when citation values are close to zero. As such we report our dependent variable in unlogged form. As we discuss in the mechanism section, there was also insufficient data to do a media analysis of the impact of *Irish Times* data.

4.1. Impact of Wikipedia on Court Judgments

If our experiment had an effect on judicial reasoning, we would expect the cases summarized on Wikipedia to be referenced more often in judicial decisions. This is shown in model-free graphical form in Figure 3. Pretreatment, the accrual of citations is highly balanced between the treatment and control groups. Posttreatment, Supreme Court cases in the treatment group are cited more frequently. To estimate the magnitude of this effect, we use the following linear regression model:

$$\begin{aligned} CaseCites_{ijt} &= \beta_0 + \beta_1 Wiki_{ij} \times After_t + \beta_2 Wiki_{ij} + \beta_3 After_t \\ &+ \pi_j + \epsilon_{ijt}, \end{aligned} \tag{1}$$

where *i* indexes cases, *j* indexes stratification blocks, and *t* indexes months. Thus, *CaseCites_{ijt}* is the number of citations to the Supreme Court case *i*, from randomization block *j*, by judicial decisions in month *t*. We construct our measure as a difference in differences, and thus our coefficient of interest is on *Wiki* × *After*, which measures the relative increase in citations to the treated group in the post period. We also include block fixed effects, π_j . Because of the stratification in our experiment, conventional standard errors will have the wrong coverage. To be more specific, after pairing cases and implementing stratification, cases of large sample imbalance become less probable or impossible. To correctly adjust standard errors for this better balance between treatment and

 Table 2.
 Balance Table on Supreme Court Cases Between Treated and Control Groups

	Variables	Average treated	Average control	t statistic
Supreme Court cases	No. of positive citations	2.3	3.6	1.30
1	No. of neutral citations	19.4	24.6	1.27
	No. of negative citations	0.30	0.55	1.79
	No. of Irish Times references	0.45	0.40	-0.35
	No. of RTE references	0.09	0.08	-0.29
	No. of other public media	0.47	0.57	1.29
	Judgment year	2009	2009	0.00
	Publication year	2009	2009	0.00
Wikipedia articles	Word count	987	978	-0.16
1	No. of external links	0.92	0.83	-0.45
	No. of academic references	3.8	3.8	0.03
	No. of nonacademic references	3.6	3.8	0.42

Note. The law category of Supreme Court cases is exactly the same between the treated and control cases within each block.



Figure 3. (Color online) Average Number of Citations per Month Before and After Treatment

control, we calculate the standard errors using randomization inference (Gadbury 2001).

The estimation results are shown in Table 3, where column (1) shows the results without block fixed effects, and column (2) those with block fixed effects. Both columns indicate that there is an increase in the number of follow-on citations of the treated group after the treatment compared with the control group and that this effect is statistically significant at the 12% and 8% level, respectively. On average, adding case summaries to Wikipedia increases monthly citations by 0.12 compared with the control group. Compared with the average number of citations per month to cases in the year prior to the treatment, 0.52, this represents an increase in citations of 23%. Although this effect is large, the statistical

Table 3. Impact of Wikipedia on the Number of Follow-on

 Citations (OLS)

	(1)	(2)
Wiki \times After	0.123	0.122*
, ,	[0.117]	[0.074]
Wiki	0.005	0.006
	[0.736]	[0.658]
After	0.085	-0.070*
2	[0.930]	[0.051]
Block fixed effects	No	Yes
No. of observations	21,126	21,126
No. of cases	154	154

Note. Randomization inference *p* values are reported in brackets. *p < 0.1; **p < 0.05; ***p < 0.01.

significance is marginal, perhaps suggesting that we are observing a heterogeneous effect.

4.2. Differences in the Impact of Wikipedia on Different Courts

In this analysis, we consider whether the effects of Wikipedia articles differ based on the type of court adjudicating the case, ranging from the High Court, which is usually the court of first instance (i.e., first deals with cases) and has high workloads, to the Court of Appeal and Supreme Court, which review appeals of lower court judgements and have lower workloads.

We might expect different responses from these different courts, as there is ample evidence that individuals tend to work faster when faced with higher workloads (Green et al. 2013, Kim et al. 2015, Kuntz et al. 2015, Berry Jaeker and Tucker 2017). Time pressure has also been shown to cause decision makers to adopt simplifying strategies (Wright 1974), such as scrolling through fewer pages on search engines (Crescenzi et al. 2021). In a legal context, we might expect these effects to manifest as a greater usage of Wikipedia by those judges, judicial assistants (clerks), and lawyers with the greatest workload.

Among Irish courts, the workload is greatest in the High Court, followed by the Court of Appeal and the Supreme Court. Of these three courts, the High Court is the only court of first instance, that is, of original jurisdiction. As such, it must first establish the facts in dispute (who did/said what) before applying the relevant law to

1954

them. Establishing the facts means taking evidence in the form of witness testimony, including cross-examination. This is not necessary in the appellate courts that, focusing on the application of the applicable law, assume that the facts are as they were originally established by the High Court (Byrne et al. 2021, p. 391). Consequently, for a High Court judge, the production of a judgment is a larger undertaking than for an appellate judge. Moreover, the greater workload per judgment of the High Court is exacerbated by the lesser support it receives from judicial assistants: Between January 2019 and January 2021, the High Court had one assistant to every 76 judgments, whereas the appellate courts had one assistant to every 26 judgments.¹⁰ The High Court's workload was specifically mentioned by Ms Justice Mary Irvine, the President of the High Court, on her recent retirement. She argued that 17 or 18 additional judges were necessary to allow the High Court to discharge its duties effectively (a 40% increase).¹¹

Hence, if time pressure was pushing judges toward Wikipedia, we would expect the strongest effects to manifest at the High Court level. Figure 4 provides model-free evidence of how the number of citations in each type of court changed in each month, affirming that all the effects seem to be arising in the High Court.

Regression results in Table 4 quantify these effects.¹² Summarizing cases on Wikipedia has a significant positive influence on the number of follow-on citations in subsequent High Court cases, but not in subsequent Supreme Court or Court of Appeal cases. Specifically, adding cases to Wikipedia increases citations in High Court cases by 0.129 per month (statistically significant at the 0.1% level). $^{\rm f3}$ Compared with 0.193 average number of citations per month to cases in the year prior to the treatment in the High Court, Wikipedia increases the subsequent citations by 66% to High Court cases. By contrast, adding cases to Wikipedia has only negligible and statistically insignificant changes to citations by the Court of Appeal or Supreme Court. Pooling these regressions allows us to test whether these differences are themselves statistically significant, and they are (regression in online appendix). These heterogeneous effects may indicate that UGC, such as Wikipedia, are more likely to be used as references in settings with high workload and time pressure.

Separate to our estimation effect, we also see an overall drop in citations occurring after the intervention, particularly by the High Court. This is likely caused by two effects. First, High Court citations have seasonal peaks and troughs every 12 months, driven by the courts' vacation schedule (OECD 2023). This can be seen clearly in Figure 3. In addition, the High Court was likely more affected by COVID because it is the court that handles most witness testimony and fact finding. Although important for an overall decrease in citations, we find that the onset of COVID is not responsible for our treatment effect (as shown later).

4.3. Dose-Response

In theorizing Wikipedia's influence on citation behavior, one might expect that greater exposure would correlate with greater impact. This is what we see. Among the treated articles, those that get more page views are associated with larger changes in citations (Table 5).

4.4. Spillovers

Having shown the direct influence of Wikipedia articles, one might wonder whether there are indirect spillover effects. For example, might a Wikipedia article link to other similar cases which then also get cited more?

If substantial, such spillover effects could potentially be a problem for our identification strategy because it would imply a stable unit treatment values assumption (SUTVA) violation as treatment write-ups drive citations to control cases. Reassuringly, a manual spot check of 26 randomly selected blocks (where the cases would be most similar) shows no evidence of the cross citations that would drive such SUTVA violations.

Although failing to find evidence of spillovers within our sample is good news for our identification strategy, it is bad news that we cannot measure the spillover effects that should exist as Wikipedia case summaries become better populated (currently, the coverage of Irish Supreme Court decisions is nearly all from our intervention). We hope that future work, particularly done in areas where there are many more Wikipedia Law articles (e.g. in the United States), chooses to investigate this interesting question.

4.5. Robustness

4.5.1. Count Models. Because our outcome variables are counts, we consider an alternative specification to our main ordinary least squares model. We have no reason to believe that the variance of our data should be equal to its mean, so we forgo using a Poisson model and instead use a negative binomial model. After using this new estimation strategy, our results shown in Table 6 remain directionally consistent and statistically significant at 1% level, providing confidence in the robustness of our findings. The incidence rate ratios (IRRs) are shown in column 1 and column 3. They show coefficients on $Wiki \times After$ of 2.070 and 1.969, respectively, depending on whether block fixed effects are used. These indicate that Wikipedia has a multiplicative effect on citations and that the magnitude of this effect is roughly a doubling. However, this effect seems to be heterogeneous. The marginal effect estimates in columns 2 and 4 show that (at the mean) adding cases to Wikipedia increases citations in subsequent cases by 11.7% or 11.2%, respectively.

4.5.2. Citation Trends and COVID-19. One might also wonder whether the COVID-19 pandemic, which



Figure 4. (Color online) Effect of Wikipedia on Different Types of Courts

occurred during our sample period, has any effect on our estimates. In particular, one might imagine that courts (or even court reporting agencies) might have behaved differently during the pandemic. Visually, this does not seem likely as the control cases citations in Figure 3 do not show any obvious trend break. Nevertheless, we empirically analyze this possibility by rerunning the analysis in Table 3, but adding a dummy

	High Court		Court o	Court of Appeal		Supreme Court	
	(1)	(2)	(3)	(4)	(5)	(6)	
Wiki \times After	0.131***	0.129***	-0.003	-0.003	-0.005	-0.005	
Wiki	-0.005	-0.003	-0.015***	-0.015***	0.024**	0.025**	
After	[0.570] 0.089***	[0.711] -0.130***	[0.000] 0.083	[0.000] 0.074	[0.019] 0.091	[0.016] -0.014	
	[0.000]	[0.000]	[0.483]	[0.482]	[0.447]	[0.683]	
Block fixed effects	No	Yes	No	Yes	No	Yes	
No. of cases	154	154	154	154	154	21,126 154	

Table 4. Impact of Wikipedia on the Number of Follow-on Citations Across Court Types (OLS)

Note. Randomization inference p values are reported in brackets.

p < 0.1; p < 0.05; p < 0.01.

variable for whether the time period in question was during the pandemic. Table 7 shows these results.

As this shows, both the magnitude and significance level are the same as those in Table 3, reflecting that although COVID-19 reduced overall citation frequencies, it does not bias our estimates. For robustness, one can also run a regression that interacts $Wiki \times After$ with *COVID* to see if the effect arises there, rather than just with $Wiki \times After$. It does not, indicating the effect has its onset with Wikification not with COVID.

4.5.3. Sensitivity Check. To check how sensitive our results are to outliers (e.g., specific cases with big changes in citations), we cycle through each block in our data, testing whether its exclusion changes our result (e.g., using leave-one-out). Among the 77 regressions that result, all remain positive and statistically significant at 1% level. Hence, no single case (or pair of them) is driving the result.

5. Mechanism Analysis

Having shown the causal link between the creation of a Wikipedia article and an increase in citations at the High Court, we consider the potential mechanisms by which Wikipedia's influence is occurring. There

Table 5. Relationship Between Wikipedia Page Views and the Number of Follow-on Citations (OLS)

	Number o	Number of citations		
	(1)	(2)		
No. of Pageviews	0.055**	0.040*		
Block fixed effects	(0.024) No	(0.021) Yes		
N	1,316	1,316		

Notes. Independent variable is log(pageview + 1). Robust standard errors are in parentheses. Regression using subsample in the posttreatment period only.

p < 0.1; p < 0.05; p < 0.01.

are many potential candidates. Judges and their staff might be using Wikipedia either directly or via a search engine that surfaces this same material. They might also be looking up the case on Wikipedia but then using the link to go to the original case (i.e., where Wikipedia is being used as a byway to the relevant legal text). Alternatively, the effect could be operating through the court filings of the lawyers in the case, who themselves use Wikipedia. In this case, judges would be getting the Wikipedia content indirectly. Equally, the effect could be operating through some third, out-of-court channel, of which there are many possibilities. For example, it could be that journalists use Wikipedia when they write up news articles and that judges or lawyers then use these media reports.¹⁴

Importantly, one should not consider this question as either/or. It is plausible that all these pathways could be contributing to the net effects that we observe. In the analysis that follows, we consider what the data can tell us about which mechanisms are at work or are likely to be. Unfortunately, because some data sources are not available, for example, lawyers' court filings, we are not able to distinguish all these effects.

5.1. Is Wikipedia Being Used as a Source of Content or a Pathway to the Underlying Case?

In one possible scenario, Wikipedia is not being used as the content source per se but just as a pathway to the underlying case, perhaps analogously to using a search engine to get to a web page. If the entire effect arose from this pathway, it would alleviate the worry that judgments depend on the content of the relevant Wikipedia articles (it would instead imply a different problem: that precedents not yet summarized on Wikipedia were being overlooked). To investigate this question, we conduct a linguistic analysis as of December 2020. Following Thompson and Hanley (2018), we analyze the linguistic similarities between Wikipedia and the

	(1) IRR	(2) Marginal effect	(3) IRR	(4) Marginal effect
Wiki × After	2.070***	0.117***	1.969***	0.112***
	(0.375)	(0.075)	(0.326)	(0.037)
Wiki	0.978	-0.005	0.936*	-0.008*
	(0.040)	(0.009)	(0.035)	(0.004)
After	0.580***	-0.090***	0.491***	-0.063***
	(0.079)	(0.018)	(0.067)	(0.009)
Block fixed effects	No	No	Yes	Yes
No. of observations	21,126	21,126	21,126	21,126
No. of cases	154	154	154	154

Table 6. Impact of Wikipedia on the Number of Follow-on Citations in High Court (Negative Binomial Model)

Note. Robust standard errors are reported in parentheses.

p < 0.1; p < 0.05; p < 0.01.

documents potentially being influenced; in our case, the subsequent citing judgments. Because we wrote summaries for the cases in both the treatment and control groups, we can look for the linguistic fingerprints of our Wikipedia articles in the judgments. If the similarities of the treatment case summaries are stronger than those of control case summaries, it would indicate that Wikipedia is influencing the linguistic content in those decisions.

To look for linguistic similarity, we apply a bag-ofwords model to create a *K*-dimensional vector where each entry denotes the frequency of a certain word to represent each document. We also use term frequencyinverse document frequency (tf-idf) weighting to account for the fact that some words are more structural, for example, "the," and thus should be downweighted, whereas others have more semantic content (Shahmirzadi et al. 2019). We then use cosine similarity (Rahutomo et al. 2012) to measure the closeness between judgments and our Wikipedia summaries, as shown in Table 8. We find a notable increase in the cosine similarity that is attributable to the Wikipedia articles, suggesting that the Wikipedia article content is itself being used.¹⁵

Table 7. COVID Effects on the Impact of Wikipedia onFollow-on Citations (OLS)

	(1)	(2)
Wiki \times After	0.123	0.122*
2	[0.117]	[0.074]
Wiki	0.005	0.006
	[0.734]	[0.659]
After	0.043	-0.020
	[0.930]	[0.758]
COVID	0.050	-0.059
	[0.495]	[0.539]
Block fixed effects	No	Yes
No. of observations	21,126	21,126
No. of cases	154	154

Note. Randomization inference p values are reported in brackets. $^{\ast}p < 0.1; \,^{\ast\ast}p < 0.05; \,^{\ast\ast\ast}p < 0.01.$ Which part of the Wikipedia article is being used also matters. If it is just the direct quotations from the cases that are put into Wikipedia, that would have different implications than if it is the writer's analysis that is being echoed. To test this, we re-estimate our results using the Wikipedia articles without the quotations. The coefficients on $Wiki \times After$ show the increase in similarity between the judgments and Wikipedia articles for cases in the treated group irrespective of whether the quotations are included in the analysis (with significance levels ranging from 3.5%–13.7%).

These results make it clear that Wikipedia is not just being used as a pathway to the underlying case. The Wikipedia content that contextualizes the case is itself influencing the language of the judgment.

5.2. Are Judges Using Wikipedia?

Ideally, we would be able to identify whether the effects we observe are due to judges and their staffs, lawyerly use, or some other cause. Of particular interest is whether judges or lawyers are themselves using Wikipedia. At least some use seems probable, given the linguistic similarity between the Wikipedia articles and the judgments. However, alternative explanations are

 Table 8. Textual Similarity Analysis (OLS)

	Full	text	Text without quo		
	(1)	(2)	(3)	(4)	
Wiki × After	0.015*	0.009*	0.014	0.012**	
Wiki	(0.009) 0.008**	(0.004) -0.009**	(0.009) 0.015***	(0.006) -0.003	
After	(0.004) -0.020**	(0.003) 0.005*	(0.004) -0.014*	(0.003) 0.006*	
	(0.007)	(0.003)	(0.007)	(0.003)	
N	No 2,811	Yes 2,811	No 2,811	Yes 2,811	

Note. Robust standard errors are reported in parentheses. *p < 0.1; **p < 0.05; **p < 0.01.

	Pos	Positive		Neutral		Negative	
	(1)	(2)	(3)	(4)	(5)	(6)	
Wiki × After	0.017**	0.017**	0.109***	0.108***	0.001*	0.001	
	[0.045]	[0.043]	[0.000]	[0.000]	[0.066]	[0.111]	
Wiki	-0.002	-0.002	-0.011	-0.009	-0.001^{*}	-0.001°	
	[0.328]	[0.348]	[0.142]	[0.181]	[0.066]	[0.093]	
After	-0.016**	-0.019**	-0.059***	-0.098***	-0.002**	-0.003*	
	[0.024]	[0.023]	[0.000]	[0.000]	[0.039]	[0.067]	
Block fixed effects	No	Yes	No	Yes	No	Yes	
No. of observations	21,126	21,126	21,126	21,126	21,126	21,126	

154

Table 9. Impact of Wikipedia on the Number of High Court Follow-on Citations Across Reference Types (OLS)

154

Note. Randomization inference p values are reported in brackets.

154

p < 0.1; p < 0.05; p < 0.01.

No. of cases

also possible. For example, it could be that the effects come through the media or alternate sources, and these are then used by lawyers or judges; although in such cases, the raw effect would need to be much *bigger* because we would be observing a diluted version as Wikipedia content is passed through intermediaries "broken-telephone" style. Unfortunately, the best data for analyzing this question, the legal filings by the lawyers, is not readily available, and thus we can only use indirect indications to infer the mechanism.

To provide additional empirical evidence on this question of lawyerly/judicial use, we examine the type of citations being made. Citations in judicial decisions can be disaggregated into three types: positive citations, where the judge invokes the case as a precedent for their decision; negative citations, where the judge distinguishes the case as not a precedent for their decision; and neutral citations, where the judge introduces the case to describe the general legal context of their decision.¹⁶ Because the valence of a citation depends on the judge's decision, equal use of Wikipedia by lawyers might be expected to produce symmetric effects. By contrast, an asymmetry might reflect that Wikipedia is being used to marshal support during the judgment-writing process, that is, by the judge or their staff. Table 9 shows that changes in citations are indeed asymmetric. After treatment, there is a large jump in positive citations¹⁷ but not in negative citations.

To understand the implications of this result, suppose that just the parties' lawyers were using Wikipedia, and that judges were merely passively influenced by the filings made to the court. In this account, Wikipedia's effect on judgments would be due solely to judges' responsiveness to the lawyers' conflicting sets of appeals to Wiki-summarized decisions. In that case, we would expect a greater prevalence of Wikisummarized judgments both among judgments cited in support of judges' preferred line of reasoning and among those cited as points of contrast. However, a

seemingly unrelated regression test (He et al. 2020, Atasoy et al. 2021) rejects the null hypothesis that the coefficients on positive and negative citations are equal (chi-square = 3.01, p = 0.0830), confirming that the effect of Wikipedia is indeed asymmetric. This indicates that the direction of the judicial ruling matters for Wikipedia summaries' impacts.

154

154

As with the linguistic similarity analysis, this result is not determinative of judicial use. Less parsimonious possibilities could also explain this pattern, for example, some combination of lawyerly use and asymmetric reliance by judges on lawyerly filings.

A third piece of data provides further evidence of asymmetry, strengthening the argument that effects are arising from within judges' offices. In Table 10, we repeat the analysis from Table 9, but consider the effects of additional Wikipedia page views on citations. Again, we find evidence that positive citations are affected, whereas negative citations are not. In this context, that means that additional searches on Wikipedia do not generate more citations that contrast with the direction of the judges' rulings but do generate more citations that support it. This is again suggestive of selective searches that favors information that aligns with the judge's preferred line of reasoning.

Table 10. Relationship Between Wikipedia Page Views and the Number of Follow-on High Court Citations Across Different Reference Types (OLS)

	Positive		Neutral		Negative	
	(1)	(2)	(3)	(4)	(5)	(6)
No. of Pageviews	0.005*	0.004	0.038***	0.031***	0.000	0.000
Block fixed effects N	No 1,316	Yes 1,316	No 1,316	Yes 1,316	No 1,316	Yes 1,316

Notes. Independent variable is log(pageview + 1). Robust standard errors are reported in parentheses. Regression using subsample in the posttreatment period only.

p < 0.1; p < 0.05; p < 0.01.

154

Taken together, the asymmetries in citation behavior combined with the linguistic echoing in judgments suggest that at least some of Wikipedia's effect is happening directly via the judges or their judicial staff.

We encourage future work to continue to explore the mechanisms behind these Wikipedia effects, particularly by trying to get access to court legal filings to establish how much effect comes through lawyers and by looking for causal effects on judicial decision making that arise through the media. (In the Irish context, there was insufficient coverage for us to have statistical power for this question.)

6. Discussion and Contributions 6.1. Wikipedia and the Rule of Law

An important theoretical contribution of our paper is to extend the impact of user-generated content into law, where legitimacy in decision making is central to ideas of fairness and to the justification of the State's coercive power (Raz 2019). This is also a methodological contribution, as we are (to the best of our knowledge) the first to run a randomized control trial on any nonprocedural aspect of judicial decision making.

Previous literature has shown the ubiquitous impact of UGC on individuals' nonprofessional decision making (Xu and Zhang 2013, Hinnosaar et al. 2023). A few studies have presented evidence that such impact of UGC is extended to professional practice, such as the software development community (Huang et al. 2022) and scientists (Thompson and Hanley 2018). Unlike prior research, our paper explores whether UGC also plays a role in the context of law, a setting where decision-making process is highly regulated and with potentially far-reaching consequences for society.

Establishing Wikipedia's influence on law has its own specific importance because the ideal of the Rule of Law promises that the state will exercise its authority to protect citizen's legal rights (Fuller 1964). Therefore, when citizens come to a court, rather than, say, a citizens' advice bureau, they come not for guidance but to vindicate their right to take (or not to take) some action. A court, in explaining how someone was entitled to do as they did, implies that all supporting legal information was supplied by competent sources: It is "... an important part of the Rule of Law that there be a competent profession available to offer ... advice ... [as to] what the law at any given time requires" (Waldron 2020; similarly, Hamilton 1788, p. 418).

Put another way, a court of law has an obligation to show how the law demands the decided outcome: "[T]he legitimacy of adjudication depends on affording those who are to be bound a right ... to be informed of the reasons for a decision" (Solum 2004, pp. 279–280). Because the expertise used to create a Wikipedia article's content is not known, a judicial decision based on Wikipedia content cannot guarantee that that the law and precedents have been expertly applied to the litigant's case. Accordingly, our finding reveals a new source of risk to traditional Rule of Law values.

Considering how to handle this challenge to Rule of Law values is also important because it previews the coming magnification of this risk by new technologies. Commercial legal databases, such as vLex, are now creating judgment summaries using large language models (vLex 2023). Like UGC, these large language models may often be right, but sometimes they can be very wrong (Choi and Schwarcz 2023, p. 4; Weiser and Schweber 2023).

6.2. Limits of Wikipedia's Influence

Our study also contributes to the literature by helping delineate the limits to UGC. By now, it is clear that UGC and other easily available sources of information have important impacts in many areas of society (Phillips et al. 1991, Ahn et al. 2016, Thompson and Hanley 2018). Although not shown causally, there are also observational studies showing that UGC has impacts in some high-stakes areas of professional practice, such as medicine (Hughes et al. 2009). From such work, it might be tempting to conclude that Wikipedia and UGC should have influence everywhere.

An opposite view is that we should expect sharp limits to Wikipedia's influence: The supports for a bridge should not be designed with Wikipedia's assessment of the carrying load of steel but instead based on the value from scientific testing. There has long been criticism of the use of UGC in professional practice. In the media industry, Hermida and Thurman (2008) criticize the quality of the content provided by amateurs despite the dramatic increase in the adoption of UGC by mainstream news organizations. Similarly, professionals in health and medicine also make use of Wikipedia, particularly medical students and junior doctors (Heilman et al. 2011, Matheson and Matheson-Monnet 2017, Smith 2020). A small-scale survey confirmed this noting that 89% of surveyed physicians used at least one Internet resource in their medical practice, with 80% using wikis, such as Wikipedia (Hughes et al. 2009). Many healthcare professional organizations and institutions have issued guidelines to regulate the use of UGC due to the potential risks caused by poor-quality information (Farnan et al. 2013, O'Hara et al. 2013, Ventola 2014). In the context of law, Hildebrandt (2016, p. 2) also states that "Technological infrastructures matter, require our attention and must somehow be brought under the Rule of Law." With such strong institutional and normative tools to limit influence, one might conclude that Wikipedia's influence should have sharp limitations in consequential decision making.

Thus, although legal professionals have embraced online information since around the turn of the millennium (Schauer and Wise 2000), legal scholars and judges have often emphasized the pitfalls of reliance on Internet sources generally, for example, "any evidence procured off the Internet is adequate for almost nothing,"¹⁸ and of UGC in particular, "Sensibly, counsel for the applicant indicated that he was not relying on these passages from the affidavit [those citing Wikipedia entries]."19 Over time, however, a more nuanced critique of UGC has emerged that focuses not its use per se but rather on the way in which it is used. In reviewing relevant case law, for instance, lawyers have been said to legitimately use UGC as long as it serves as a supplemental and not a standalone research source (Novak 2010). Equally, it is notable that both scholars and judges have characterized lawyers' reliance on free Internet content as a question of resources, explaining its employment as a means to help "save litigants time and money" (Whiteman 2010).²⁰

Faced with conflicting views about the extent of UGC's influence, it is valuable to look for empirical answers to sharpen our understanding of those boundaries. Our results show that normative prohibitions do seem to keep Wikipedia's influence out of the most significant, well-resourced parts of law, as represented by the Irish Court of Appeals and Supreme Court, where we observe no impact from the addition of our Wikipedia articles. By contrast, we see that these prohibitions are insufficient for keeping Wikipedia's influence out of High Court decisions, where time and resources pressures would incentivize the use of Wikipedia.

Hence, our research helps make clear the limits of UGC's influence. It suggests that normative mechanisms can be put in place to curb Wikipedia's reach but that such measures should be seen as disincentives that may be counterbalanced by incentives that favor use, such as ease of access.

6.3. Importance of Reliable UGC

In addition to contributing to the normative debate about where UGC should be used, our paper also contributes to the practical debate about how to manage organizations doing knowledge production and how policy solutions can improve these outcomes. There exists a broad debate on the accuracy and reliability of UGC (Kane and Ransbotham 2016, Greenstein and Zhu 2018).

Faced with such concerns, managers in knowledge production should be worried about our results. This could be in law, with judges wanting to stop their staff from using Wikipedia or lawyers wanting to stop legal assistants from using it. It could also be outside of law, for example, editors wanting to stop reporters from using Wikipedia. Our results reveal that even in fields that prize authoritative knowledge, the ubiquity and visibility of Wikipedia make it a temptingly easy reference source for professionals doing their work. Managers supervising employees in this situation would be well to heed this and consider work practices that are cognizant of these effects. For example, firms could embrace this use but assign time from experts in the company to ensure that the Wikipedia content does not contain any misleading information. Similarly, they could institute double checks for important claims, where fact checkers go back to primary sources. Alternatively, the firm could engage in other practices to limit the influence of Wikipedia, for example, website blocking, instituting norms against its use, and so on. That said, it is unclear how well these latter methods would work because they have not worked well in areas like academia.

Perhaps even more important than responses by individual managers, our results speak to the potential for policy that provides a rich informational infrastructure in the way that it already provides physical infrastructure. In particular, it suggests that experts (say the bar association or the courts) could be funded to improve the quality of content on Wikipedia (or similar repositories) and that such interventions could have broad impact on improving knowledge production. Just as drivers can use a high-quality road to travel better, knowledge workers could use high-quality information to work better.

7. Conclusion

UGC is an established feature of the digital landscape, and Wikipedia is perhaps the most notable platform in it, providing content to millions of information seekers every day. We investigate whether Wikipedia affects legal decision making by conducting a randomized control trial in which we add Wikipedia articles on Irish Supreme Court precedents. Our study extends the existing knowledge of UGC by showing the surprising reach of Wikipedia even into highly formal settings where processes are tightly prescribed.

We find that summarizing Irish Supreme Court cases on Wikipedia increases their citation in subsequent legal judgments. This effect does not extend to higher courts (Court of Appeals, Supreme Court) but rather is confined to courts that decide cases at first instance, where workload pressure is greater and the temptation to take advantage of Wikipedia's easy access may be stronger. This distinction is theoretically meaningful because it helps articulate the boundaries of UGC influence.

We also explored the mechanisms underpinning Wikipedia's influence. Linguistic similarities in the judgment text indicate that the articles themselves are providing context and that they are relied on directly in framing legal arguments. Asymmetries in both the experiment's effects on citation behavior and in the correlations between Wikipedia page views and citation types also suggest that judges (or their staffs) are themselves using Wikipedia (although other, less parsimonious explanations are also possible).

By showing that Wikipedia is influencing the practice of law, our article provides further evidence of the importance of having informative, unbiased content in these important public repositories of knowledge.

Acknowledgments

The authors thank Michael Doherty and Maria Murphy for valuable contributions, Jana Khalil for research assistance, and the students of Maynooth University's School of Law and Criminology who assisted with Wikipedia article drafting.

Endnotes

¹ See, for example, *Fire Insurance Exchange v. Oltmanns* (2012) Case No. 20100462-CA-Utah Court of Appeals, where the court referred to Wikipedia when discussing the use of the term "jet ski."

² Below the superior courts, there are also the Circuit Court and District Court, which are courts of first instance with limited jurisdiction.

³ These structural and ideological similarities between the Irish and other common-law judicial systems have resulted in the Irish courts taking guidance from other jurisdictions, including the United States. In the Irish case of *State (Woods) v. Attorney General* [1969] IR 385 (Ir.) for example, Justice Henchy used a rule of construction that" accords with the practice of the American Supreme Court: *US v Delaware & Hudson Co; US v Witkovich*" when considering the constitutionality of a statute.

⁴ The Courts Service of Ireland, https://www.courts.ie/judgments, accessed December 28, 2021.

⁵ See the Irish Courts' Practice Directions https://www.courts.ie/ content/practice-directions.

⁶ This paper covers an experiment also discussed in (Thompson et al. 2024). That paper focuses on the legal implications for judicial practice and the empirical contribution to studies of jurisprudence. This paper differs both in terms of focus (on UGC) and by having a much greater depth of empirical analysis). All data and materials for this paper are available on the Open Science Framework. https://osf.io/mytqf/?view_only=34fcff72f47f4a6cbe3c45cb558e7e32.

⁷ JustisOne is an online, subscription-based legal information platform (now replaced by vLex Justis) that aggregates judgments of (among others) the Irish courts from the official Irish Courts Service website, as well as law report versions where available. This database provides case information on the relevant category of law covered in the case, the number of subsequent citations, court types, and reference types. We also cross-checked these against the courts' official website (courts.ie) and found less than a 1% difference.

⁸ See, for instance, our Wikipedia entry for the decision in *Weir-Rodgers v. SF Trust Ltd.* https://en.wikipedia.org/wiki/Geraldine_Weir-Rodgers_v._SF_Trust_Ltd. A screenshot of the Wikipedia page is in the online appendix.

⁹ AsPredicted #70696 https://aspredicted.org/download_pdf.php?b= cyT1zTihHlsPqLplqDIHzBP5Gl2zKslKGgUZWxD5N6S7k9nEn8&a=dk 13V1RwaWN1VIJyY0gwaU56ZFVOUT09.

¹⁰ Judgment information retrieved from the Irish Courts Service www.courts.ie. Information on judicial assistants derived from a composite of sources: Irish Courts Annual Report (2019), OECD (2023), and Irish Supreme Court Annual Report (2020).

¹¹ Phelan (2022), "'Lack of judges is hampering justice,' claims retiring High Court President Mary Irvine." Justice Irvine's comments are consistent with those of Mr Justice David Barniville, the current President of the High Court, who has called for an additional 20 High Court judges (Carolan 2022).

¹² We conducted a formal regression using an event study specification (regression shown in the online appendix) to show that there is no significant influence of Wikipedia on citations in the pretreatment period. Therefore, we see no evidence that the parallel trend assumption has been violated.

¹³ Given that citations present strong temporary patterns in the model-free evidence, we rerun the regression using calendar month fixed effects (regression shown in the online appendix). The results are similar in magnitude and significant level. These results are also robust to using a balanced panel, focusing on months -28 to +7 and clustering standard errors at the case level (see corresponding regression in the online appendix), and to using an event-study construction instead of a simple pre-vs-post structure.

¹⁴ We tested the media as a mechanism using the U.S. Supreme Court case and the New York Times data because the Irish Times did not have enough data points to conduct the analysis. The results are shown in Online Appendix A5.

¹⁵ Unfortunately, despite being a workhorse technique in natural language processing, cosine similarity is more useful for detecting effects and ordinal ranking but resists cardinal interpretation (Ehrmann and Talmi 2020, Girardi et al. 2021).

¹⁶ JustisOne (now vLex Justis) does this categorization. There are also 9% of citations that they have left uncategorized, which we drop from our analysis.

¹⁷ Although the coefficient on positive citations might appear small, it nevertheless represents a notable change of 85%.

¹⁸ Cant Samuel 1999 St. Clair v. Johnny's Oyster & Shrimp, Inc., 76 F. Supp. 2d 773, (S.D. Tex.)

¹⁹ Birmingham George 2012 Rowan v Kerry County Council (High Court of Ireland) [IEHC] 65. para 31; similarly, Margolis (2007).

²⁰ See also Martinez William Lyall v. City of Denver, 319 F.R.D. 558, 569 (D. Colo. 2017).

References

- Ahn DY, Duan JA, Mela CF (2016) Managing user-generated content: A dynamic rational expectations equilibrium approach. *Market-ing Sci.* 35(2):284–303.
- Atasoy H, Banker RD, Pavlou PA (2021) Information technology skills and labor market outcomes for workers. *Inform. Systems Res.* 32(2):437–461.
- Berman HJ (2000) The western legal tradition in a millennial perspective: Past and future. LA Law Rev. 60(3):739–763.
- Berry Jaeker JA, Tucker AL (2017) Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity. *Management Sci.* 63(4):1042–1062.
- Byrne R, McCutcheon P, Cahillane L, Roche-Cagney E (2021) *Byrne and McCutcheon on the Irish Legal System* (Bloomsbury Professional, London).
- Carolan M (2022) More judges and extra working day for some District Court judges expected in working group's recommendations. *Irish Times* (October 9), https://www.irishtimes.com/crimelaw/courts/2022/10/09/more-judges-and-extra-working-dayfor-some-district-court-judges-expected-in-working-groupsrecommendations/.
- Chhabra A, Iyengar SRS (2020) Who writes Wikipedia? An investigation from the perspective of Ortega and Newton Hypotheses. Proc. 16th Internat. Sympos. on Open Collaboration (ACM, New York), 11.
- Choi J, Schwarcz D (2023) AI assistance in legal analysis: An empirical study. Preprint, submitted August 16, https://dx.doi.org/10. 2139/ssrn.4539836.

- Courts Service Annual Report (2020) The courts service of Ireland. Accessed February 1, 2024, https://www.courts.ie/acc/alfresco/ b47652ff-7a00-4d1f-b36d-73857505f860/Courts_Service_Annual_ Report_2020.pdf/pdf#view=fitH.
- Crescenzi A, Capra R, Choi B, Li Y (2021) Adaptation in information search and decision-making under time constraints. *Proc.* 2021 ACM SIGIR Conf. on Human Inform. Interaction and Retrieval (ACM, New York), 95–105.
- Ehrmann M, Talmi J (2020) Starting from a blank page? Semantic similarity in central bank communication and market volatility. *J. Monetary Econom.* 111:48–62.
- Engel C, Weinshall K (2020) Manna from heaven for judges: Judges' Reaction to a Quasi-Random Reduction in Caseload. J. Empirical Legal Stud. 17(4):722–751.
- Fader PS, Winer RS (2012) Introduction to the special issue on the emergence and impact of user-generated content. *Marketing Sci.* 31(3):369–371.
- Farnan JM, Sulmasy LS, Worster BK, Chaudhry HJ, Rhyne JA, Arora VM, Professionalism American College of Physicians Ethics, et al. (2013) Online medical professionalism: Patient and public relationships: Policy statement from the american college of physicians and the federation of state medical boards. *Ann. Internal Medicine* 158(8):620–627.
- Forman C, van Zeebroeck N (2019) Digital technology adoption and knowledge flows within firms: Can the Internet overcome geographic and technological distance? *Res. Policy* 48(8):103697.
- Fuller LL (1964) *The Morality of Law* (Yale University Press, New Haven, CT).
- Gadbury GL (2001) Randomization inference and bias of standard errors. Amer. Statist. 55(4):310–313.
- Gerken JL (2010) How courts use Wikipedia. J. Appellate Practice Processes 11:191.
- Girardi G, Hanley KW, Nikolova S, Pelizzon L, Sherman MG (2021) Portfolio similarity and asset liquidation in the insurance industry. J. Financial Econom. 142(1):69–96.
- Green LV, Savin S, Savva N (2013) "Nursevendor problem": Personnel staffing in the presence of endogenous absenteeism. *Management Sci.* 59(10):2237–2256.
- Greenstein S, Zhu F (2018) Do experts or crowd-based models produce more bias? Evidence from Encyclopedia Britannica and Wikipedia. *Management Inform. Systems Quart.* 42(3):945–959.
- Hamilton A (1788) The Federalist No. 78, 418, 2005 ed. (J.R. Pole).
- He S, Peng J, Li J, Xu L (2020) Impact of platform owner's entry on third-party stores. *Inform. Systems Res.* 31(4):1467–1484.
- Heilman JM, Kemmann E, Bonert M, Chatterjee A, Ragar B, Beards GM, Iberri DJ, et al. (2011) Wikipedia: A key tool for global public health promotion. J. Medical Internet Res. 13(1):e14.
- Hermida A, Thurman N (2008) A clash of cultures: The integration of user-generated content within professional journalistic frameworks at British newspaper websites. *Journalism Practice* 2(3):343–356.
- Hildebrandt M (2016) Law *as* information in the era of data-driven agency. *Modern Law Rev.* 79(1):1–30.
- Hinnosaar M, Hinnosaar T, Kummer M, Slivko O (2023) Wikipedia matters. J. Econom. Management Strategy 32(3):657–669.
- Hitt M (2019) Inconsistency and Indecision in the United States Supreme Court (University of Michigan Press, Ann Arbor, MI).
- Huang P, Ceccagnoli M, Forman C, Wu DJ (2022) It knowledge spillovers, absorptive capacity, and productivity: Evidence from enterprise software. *Inform. Systems Res.* 33(3):908–934.
- Hughes B, Joshi I, Lemonde H, Wareham J (2009) Junior physician's use of Web 2.0 for information seeking and medical education: A qualitative study. *Internat. J. Medical Inform.* 78(10): 645–655.
- Irish Courts Annual Report (2019) The courts service of Ireland. Accessed February 1, 2024, https://www.courts.ie/acc/alfresco/ 9bd89c8a-3187-44c3-a2e9-ff0855e69cb5/Courts%20Service%20 Annual%20Report%202019.pdf/pdf#view=fitH.

- Irish Supreme Court Annual Report (2020) The courts service of Ireland. Accessed February 1, 2024, http://www.supremecourt.ie/ SupremeCourt/sclibrary3.nsf/pagecurrent/A48207B9CF7E7B 46802583B00040E2EA?opendocument&l=en.
- Kane GC, Ransbotham S (2016) Research note-content and collaboration: An affiliation network approach to information quality in online peer production communities. *Inform. Systems Res.* 27(2): 424–439.
- Kim SH, Chan CW, Olivares M, Escobar G (2015) ICU admission control: An empirical study of capacity allocation and its implication for patient outcomes. *Management Sci.* 61(1):19–38.
- Kuntz L, Mennicken R, Scholtes S (2015) Stress on the ward: Evidence of safety tipping points in hospitals. *Management Sci.* 61(4): 754–771.
- Lobingier S (1946) Precedent in past and present legal systems. *Michigan Law Rev.* 44:955–966.
- Margolis E (2007) Surfin'safari-why competent lawyers should research on the web. Yale JL Tech. 10:82.
- Matheson D, Matheson-Monnet C (2017) Wikipedia as informal selfeducation for clinical decision-making in medical practice. Open Medicine J. 4(suppl 1, M2):15–25.
- Novak M (2010) Legal research in the digital age: Authentication and preservation of primary material. Accessed February 1, 2024, https://digitalcommons.unl.edu/lawlibrary/4/.
- OECD (2023) Modernising Staffing and Court Management Practices in Ireland: Toward a More Responsive and Resilient Justice System (OECD Publishing, France).
- O'Hara B, Fox BI, Donahue B (2013) Social media in pharmacy: Heeding its call, leveraging its power. J. Amer. Pharmaceutical Assoc. (Washington DC) 53(6):565.
- Okoli C, Mehdi M, Mesgari M, Nielsen FÅ, Lanamäki A (2014) Wikipedia in the eyes of its beholders: A systematic review of scholarly research on Wikipedia readers and readership. J. Assoc. Inform. Sci. Tech. 65(12):2381–2403.
- Peoples LF (2009) The citation of Wikipedia in judicial opinions. *Yale JL Tech.* 12:1.
- Pew Research Center (2016) Wikipedia at 15: Millions of readers in scores of languages. Accessed February 1, 2024, https://www. pewresearch.org/fact-tank/2016/01/14/wikipedia-at-15/#:text =Wikipedia%20averages%20more%20than%2018,to%20data%20 reported%20by%20Wikipedia.
- Phelan S (2022) Lack of judges is hampering justice, claims retiring High Court President Mary Irvine. *Irish Independent* (July 14), https://www.independent.ie/irish-news/courts/lack-of-judgesis-hampering-justice-claims-retiring-high-court-president-maryirvine-41838374.html#:text=Ms%20Justice%20Mary%20Irvine, judges"%20in%20the%20High%20Court.
- Phillips DP, Kanter EJ, Bednarczyk B, Tastad PL (1991) Importance of the lay press in the transmission of medical knowledge to the scientific community. *New England J. Medicine* 325(16):1180–1183.
- Rahutomo F, Kitasuka T, Aritsugi M (2012) Semantic cosine similarity. Proc. 7th Internat. Student Conf. on Advanced Sci. and Tech., vol. 1, 1.
- Raz J (2019) The law's own virtue. Oxford J. Legal Stud. 89(1):1.
- Schauer F, Wise VJ (2000) Nonlegal information and the delegalization of law. J. Legal Stud. 29(S1):495–515.
- Schiavoni JS (2002) Who's afraid of precedent: The debate over the precedential value of unpublished opinions. UCLA Law Rev. 49:1859.
- Sen A, Grad T, Ferreira P, Claussen J (2023) (How) does user-generated content impact content generated by professionals? Evidence from local news. *Management Sci.*, ePub ahead of print October 23, https://doi.org/10.1287/mnsc.2023.4962.
- Shahmirzadi O, Lugowski A, Younge K (2019) Text similarity in vector space models: A comparative study. Proc. 18th IEEE Internat. Conf. on Machine Learn. and Applications (IEEE, Piscataway, NJ), 659–666.
- Smith DA (2020) Situating Wikipedia as a health information resource in various contexts: A scoping review. *PLoS One* 15(2):e0228786.

Solum LB (2004) Procedural justice. Southern California Law Rev. 78:181-321.

- Thompson N, Hanley D (2018) Science is shaped by Wikipedia: Evidence from a randomized control trial. Preprint, submitted September 20, https://dx.doi.org/10.2139/ssrn.3039505.
- Thompson N, Flanagan B, Richardson E, McKenzie B, Luo X (2024) Trial by Internet: A randomized field experiment on Wikipedia's influence on judges' legal reasoning. Tobia K, ed. *The Cambridge Handbook of Experimental Jurisprudence* (Cambridge University Press, Cambridge, UK), Forthcoming.
- Ventola CL (2014) Social media and healthcare professionals: Benefits, risks, and best practices. *Pharmacy and Therapeutics* 39(7):491.
- vLex (2023) vLex blog post. Accessed February 1, 2024, https://vlex. com/blog/vincent-case-analysis.
- Waldron J (2020) The rule of law. Accessed February 1, 2024, https://plato.stanford.edu/archives/fall2023/entries/rule-of-law/.

- Weiser B, Schweber N (2023) The ChatGPT lawyer explains himself. New York Times (June 8), https://www.nytimes.com/2023/06/ 08/nyregion/lawyer-chatgpt-sanctions.html.
- Whiteman M (2010) The death of twentieth-century authority. UCLA Law Rev. Discourse 58:27.
- Wikipedia (2022) Wikipedia: Statistics. Accessed August 24, 2022, https://en.wikipedia.org/wiki/Wikipedia:Statistics.
- Wikipedia (2023) Wikipedia: General disclaimer. Accessed September 10, 2023, https://en.wikipedia.org/wiki/Wikipedia:General_ disclaimer.
- Wright P (1974) The harassed decision maker: Time pressures, distractions, and the use of evidence. J. Appl. Psych. 59(5):555.
- Xu SX, Zhang X (2013) Impact of Wikipedia on market information environment: Evidence on management disclosure and investor reaction. *Management Inform. Systems Quart.* 37(4):1043–1068.