Bad Science:
Retractions, Citations and Media Coverage.

Eleonora Alabrese*

February 18, 2022

PRELIMINARY - DO NOT CIRCULATE

Abstract

Bad science can be hard to eradicate. This creates the potential for dissemination of misinformation within and outside academia. This work shows that media coverage can be helpful to the auto-correcting process of science. I use a conditional difference-in-differences strategy to show that retracted articles experience larger citation losses in the presence of media coverage and the remaining post-citations mention more often that the paper is indeed retracted. I further show suggestive evidence of selection into treatment for papers attracting excess coverage and that journals that generally publish popular articles are those where retractions happen faster and where citation penalties are larger. The differential effect of media coverage is observed only for hard sciences, suggesting distinct publication practices may impact the visibility of the retraction. I finally show that newspapers are more likely to cover the publication of a paper rather than its retraction, an imbalance that could impact public perception of scholars' trustworthiness.

Keywords: Science; Retractions; Citations; Media Coverage; Misinformation; Altmetric.

*Acknowledgements: I thank Thiemo Fetzer, Sascha Becker and Andreas Stegmann for the precious mentoring and guidance. I want to thank Arun Advani, Wiji Arulampalam, Manuel Bagues, Ludovica Gazze, Vardges Levonyan, Andrew Oswald, Kirill Pogorelskiy, Carlo Schwarz, Natalia Zinovyeva and participants of the Workshop on Political Economy (UniBZ-Brunek), the Political Economy and Public Economics Research Away Day (Warwick) and Warwick internal seminars for insightful comments and suggestions. University of Warwick. Email address: e.alabrese@warwick.ac.uk.

.

1 Introduction

Bad science is a well studied phenomenon (see Hesselmann et al., 2017, for a review on research on scientific retractions) and the fact that fraudulent work continue to appear long after it is meant to be removed from the literature is one of the main concerns for scientists (Schneider et al., 2020). In a moment where salience around the scientific process is at its highest this aspect is attracting more and more general attention.[1] Even recently, examples of well published fraudulent Covid-19 research (Mehra et al., 2020a,b) are still heavily cited (Piller, 2021) despite being discredited in global headlines. Eminent scholars of any field could be involved in such events as for the recent case of Dan Ariely, a behavioral economist whose paper on dishonesty was recently retracted.[2]
People's failure to update their believes could explain the persistence of discredited ideas (Goncalves et al., 2021). This creates the potential for dissemination of misinformation within and outside academia (Lewandowsky et al., 2012). One question then becomes: what is the role of media reporting in the scientific process? In the scientific context, the visibility and accessibility of retractions and their associated retraction notices are decisive and yet no communication channel is flawless. This paper explores the role of mainstream media in the retraction process, a channel largely unexplored by the literature.
Understanding the impact of media coverage might be challenging given its multifaceted nature. In terms of salience, media coverage may advertise the findings of a study leading to higher future citations (Phillips et al., 1991; Fanelli, 2013; Ziegler, 2021). At the same time, media coverage could exert a monitoring role by criticizing a study and updating the public about the status of a paper (e.g. its retraction). In a related work, Whitely (1994) looked at the case of Robert A. Slutsky, an author famously known to have published fraudulent articles, and found that after his work was questioned in the news media (1985) scientists cited his publications less compared to controls. Nonetheless, the

[1](see The Economist, 2021)

[2]https://retractionwatch.com/2021/09/14/highly-criticized-paper-on-dishonesty-retracted/

prominence of the scandal and the institutional setting at that time, makes this case rather unique. More in general, Serghiou et al. (2021) reveals that retracted articles may receive high coverage but pre-retraction coverage far outweighs post-retraction coverage. Finally, in terms of selection, the work of eminent authors could be more likely to attract media coverage (e.g. Ivanova et al., 2013) but those authors are at times the ones less impacted by the consequences of retractions (Azoulay et al., 2017; Jin et al., 2019). Hence, whether media leads to more or less scrutiny and differential punishment is a priori ambigous.

This work therefore investigates whether media coverage of retracted articles could serve as an ally in the auto-correcting process of science, where articles' online coverage is directly observed across several outlets and across time.[3]

The main focus of the analysis is understanding changes in scientists' assessment of the quality of discredited work with different media exposure, and this is best proxied by variation in the citation pattern of retracted papers. To study such changes, I employ a conditional difference-in-differences strategy to examine the citation effect of a retraction shock comparing the pre-post differences for treatment papers (retracted papers) with the pre-post differences for control papers (never retracted papers). I then further compare pre-post differences across papers with diverse media coverage. However, research articles that attract online coverage might be peculiar in many ways, among which the salience of the topic under investigation, which could create concerns of selection into treatment. To solve this issue, control papers are carefully selected to mimic pre-retraction characteristics in terms of journal and year of publication, citations trends (pre-retraction) and online coverage (close to publication), such that they could best simulate the citation path of retracted papers absent the retraction.[4]

More in detail, I focus on a sample of retractions published after 2010[5] and

[3]I look at the frequency and distribution of online mentions of research articles as they appeared in newspapers, blogs, and social media. A core part of the analysis distinguishes research articles with or without news and blogs mentions within two weeks from publication. This information is collected by and extracted from *Altmetrics*.

[4]Furman et al. (2012); Lu et al. (2013); Azoulay et al. (2015); Mongeon and Larivière (2016); Azoulay et al. (2017); Jin et al. (2019) all adopt somewhat similar difference-in-differences strategies to estimate the effect of a retraction shock.

[5]Consistent data on media coverage is only recent. Hence, I select papers that are published (and retracted) strictly after 2010.

extracted from the *RetractionWatch* database, and to capture potential heterogeneity across disciplines, I select papers published in highly ranked journals across all available fields.[6] Using this sample, I first demonstrate that retracted papers do attract some media coverage, although publication and retraction events feature differently across outlets. Newspapers are more likely to cover the publication of a paper while blogs are more likely to inform about its retraction. Second, I illustrate that retractions have a negative impact on citations of retracted papers and that impact is robustly exacerbated by media coverage (30% additional reduction in forward yearly citations on average). Yet, the effect differs by outlet and coverage intensity. Third, I show suggestive evidence of a small selection into treatment (retraction) for papers that attract some media coverage at the time of publication. Predicted media coverage based on words in papers' titles is here used to control for topic endogeneity and allows to separate the differential impact of the remaining (arguably) exogenous coverage. Fourth, I show that journals that generally publish more popular articles[7] are journals where retractions happen faster (one standard deviation increase in journal visibility implies a reduction in the timing of retraction of 15%) and where citation penalties for retracted papers are larger.[8] Fifth, I uncover that the additional impact of media coverage on post-citations is evident for hard sciences as opposed to social sciences, which may suggest distinct publication practices may impact the visibility of a retraction.[9]

Based on these findings, I hypothesize that the additional effect of media could be driven by two different mechanisms: (a) higher scrutiny by the scientific community to a paper that gained publicity; (b) additional information provided to some part of the scientific community which would have otherwise remained unaware of the retraction. To investigate for the presence of the latter information channel, I check whether the textual content of post-citations significantly differs in presence of media coverage. The idea is that without the additional

[6] I retain retractions that appeared in *Scimago* top ten journals for each available discipline or in *Google* top publication journals for each available category.

[7] A journal visibility is calculated averaging the media coverage of non retracted articles published in it.

[8] This latter result also corroborates the previously mentioned media effect on yearly citations.

[9] Alternatively one can imagine that the audience may expect findings to be less "absolute" in social sciences as the object of study could be considered more volatile.

information mechanism I should observe that popular retractions experience a lower quantity of post-citations but I might not see any change in their level of "accuracy". Instead, I discover that in presence of media coverage the remaining post-citations more often mention that the paper is indeed retracted. This finding seems to confirm that in presence of media coverage scientists are more aware of a retraction and correctly acknowledge it when citing the original paper, hence reducing the potential for misinformation.

This paper contributes to several strands of the literature. First, the paper closely relates to the large literature on retractions across fields and specifically to studies that look at the causal effect of a retraction shock on citations of retracted papers (Furman et al., 2012), of authors' previous publications (Lu et al., 2013; Azoulay et al., 2017; Jin et al., 2019) and potential spillover on the field (Azoulay et al., 2015). In this paper, I use a methodology close to the ones developed in these studies which I update to incorporate the media exposure of selected papers. More in general, while a few studies in the literature on retractions mention the role of media (Sugawara et al., 2017; Sarathchandra and McCright, 2017; Serghiou et al., 2021) I here conduct a general investigation on the potential influence of media coverage in the retraction process.

The paper also relates to the literature investigating the relationship between science and the media (Weingart, 1998; Phillips et al., 1991; Fanelli, 2013; Ivanova et al., 2013; Sumner et al., 2014; Dumas-Mallet et al., 2020; Ziegler, 2021) to which I contribute by showing that media coverage of eventually retracted papers can influence reputational considerations within the scientific community.

This work also contributes to the literature on factors influencing citation rates (for example see Card and Dellavigna, 2020; Card et al., 2020; see also Tahamtan et al., 2016, for a review of the literature) to which I contribute by illustrating that the salience of a paper is a relevant factor impacting citations in case of a negative event affecting the reputation of a paper.

Finally, this work relates to the large literature on misinformation and how media channels influence politics and public policies more in general (for example see Allcott and Gentzkow, 2017; Lazer et al., 2018; see also Prat and Strömberg, 2013, for a review of the literature). I contribute to this literature by showing that on one hand, media coverage seems to attenuate the potential for misinforma-

tion within the scientific community, on the other, I illustrate that newspapers cover more the publication of a paper rather than its retraction which creates potential for disseminating misinformation to a much larger audience.
In the remainder of the paper, I examine the institutional context for retractions. I then turn to data, methods, and a detailed presentation of results.

2 Context

Bad science in not a new phenomenon and it is rather persistent. Understanding the incentives and governance regulating scientific knowledge production, dissemination and accumulation is therefore crucial to this work.
One of the most discussed institutional setting is the peer-review system. Articles are submitted and reviewed by independent experts before being accepted for publication. This feature is used to maintain high quality standards while allowing a suitable publication timing, even though practices vary greatly across disciplines and journals. This system eventually provides only limited guarantee against bad science.
Another aspect is the practice of citing related literature which is crucial for scientific communication. It allows to effectively contextualise a research article with respect to pre-existing literature while acknowledging original contributions from previous authors. Citations are regarded as an indicator of the importance of scientific findings and of their creators and can be negatively impacted by a retraction (Furman et al., 2012; Lu et al., 2013; Azoulay et al., 2015; Mongeon and Larivière, 2016; Azoulay et al., 2017; Jin et al., 2019).
In academic publishing, a retraction is the result of a procedure used by journals to alert readers that a published article should be removed from the literature. A retraction may occur when a major error (e.g. in the analysis or methods) invalidates the conclusions of the article, or in presence of misconduct (e.g. fabricated data, manipulated images, plagiarism, duplicate publication, research without required ethical approvals etc). It differs from a correction issued in case of an error or omission which can impact the interpretation of the article, but where the scholarly integrity remains intact. The surge in the absolute number of retractions across all disciplines has alarmed many in the scientific community

(see Figure 1). Nonetheless, retractions remain relatively rare involving 4 in every $10,000$ published papers of which 60% due to some type of misconduct, though both rates have been rising steadily over time (Brainard, 2018).

A retraction can be initiated by the editors of a journal, by some or all the authors or their institution and are typically complemented by a notice meant to clarify the reason of such decision. But, the information contained in notices vary significantly, some explain the details which lead to the retraction outcome and inform on whether an article results and conclusion should be disregarded entirely or in part, others are rather succinct and vague.

Figure 1: Retractions over time and across subjects.

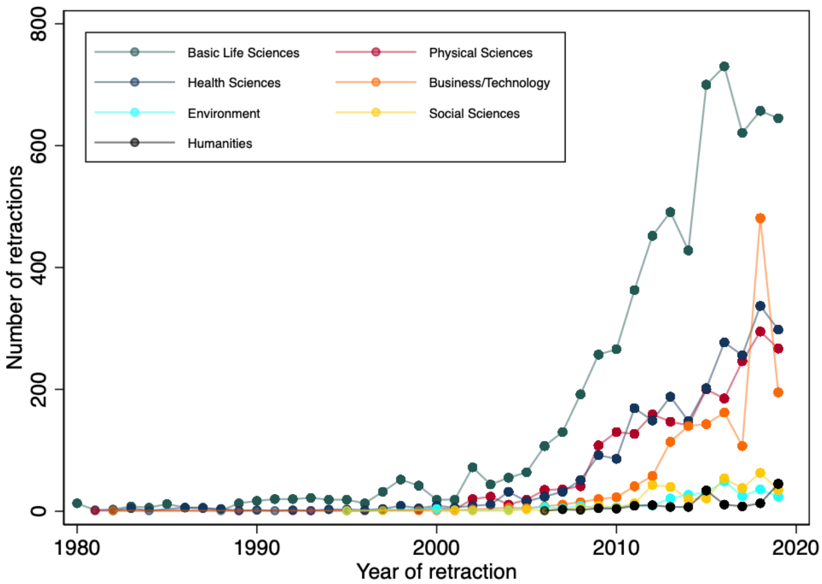

Note: Numbers reflect the full *RetractionWatch* database as of July 2020, for visual purposes one outlier publisher (e.g. IEEE) was excluded.

A further element of discussion is therefore the visibility and accessibility of both retractions and notices. "Authors are responsible for checking that none of the references cite retracted articles except in the context of referring to the retraction" (International Committee of Medical Journal Editors 2019). Awareness of readers is therefore decisive and yet no communication channel is flawless. Retractions are usually published and linked to the original publication and can be often identified via different sources (e.g. libraries, databases and search engines) but still be ignored. Schneider et al. (2020) finds that in the case of an infamous clinical trial (Matsuyama et al., 2005), in which data were falsified leading to a retraction in 2008, the retraction is not mentioned by 96% of post-retraction citations and 41% of these inaccurate citations describe the paper in detail leading to possible disinformation.

Efforts have been made to alert scientists as in the case of the specialised blog *RetractionWatch* which reports on retractions and gathers information surronding specific retraction events, such as which of the authors is responsible for the article ultimate fate. Information which is usually hard to acquire based on the notice alone. New tools are also emerging as in the case of *Scite.ai* a recently launched platform which categorises references, monitors retracted papers by searching through Crossref, PubMed, and the RetractionWatch database, and flags both citing and retracted papers on Twitter. Yet, Piller (2021) looks at the case of two high-profile Covid-19 retractions (Mehra et al., 2020a,b) on two influencial medical journals and finds that 52.5% of the citations do not correctly mention the paper status. We do not fully understand the reason why the amount of appropriate citations differs significantly in the two retraction examples above,[10] but the global interest attracted by the Covid-19 retractions might have given emphasis to the status of those papers.

Popular media, which do not necessary target the scientific community alone, like newspapers, blogs and social media accounts, have been recently active in advertising retracted articles (see Figure 2).[11] In general, media platforms seem to cover both original contributions and retractions, but the two events feature to a different extent across outlets, giving raise to potential disinformation. Indeed, Figure 3 shows that mentions in newspaper articles appear predominantly close to the publication date of a study and generally inform the public about its discovery, less often this information is updated with a new mention at the time of retraction. At the same time, mentions in blogs occur mostly around the retraction event. These blogs are often specialized and directly target academics while a wider audience is exposed to information which is not always complete, which in turn could lead to unintended consequences.

In essence, the rise of the internet and the appereance of new platforms has the

[10] Matsuyama et al. (2005) as compared to Mehra et al. (2020a,b).

[11] Notice that more recent years are likely underreported given retractions take some time to arise and hence feature in the database.

potential to direct scientists (and non-scientists) attention towards "interesting" contributions which in some cases prove to be less reliable (Serra-Garcia and Gneezy, 2021). It is therefore important to investigate whether positive post-retraction citations, and the retraction process more in general, relate to the visibility of a research paper and its retraction.

Figure 2: Media coverage of sample of retracted papers.

Panel A: Number of retractions with media

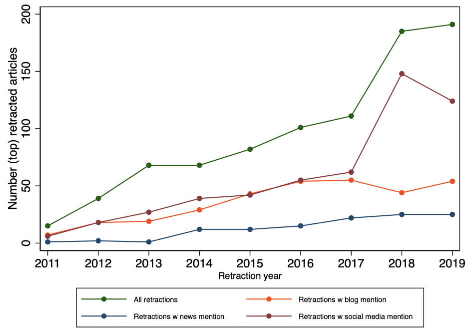

Panel B: Share of retractions with media

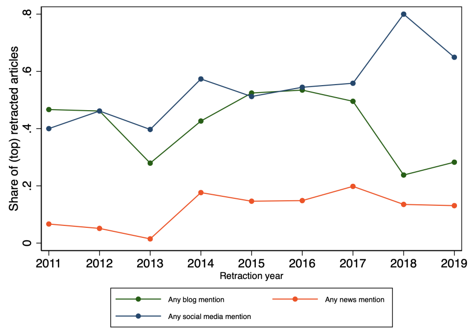

Panel C: Share of retractions with early mentions

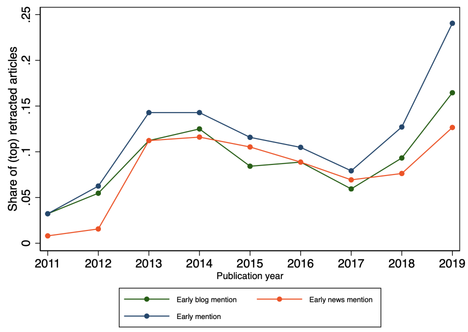

Note: Panel A shows the absolute number of retracted articles in the sample (green) which ever featured in blogs (orange), newspapers (blue), or social media (red), ordered by the year when the retraction occurred. Panel B shows the share of retracted papers that ever appeared in blogs (green), newspapers (orange), or social media (blues), again ordered by year of retraction. Panel C represents the share of retracted articles that were ever mentioned in blogs (green), newspapers (orange) or at least one of the two (blue) within two weeks from pubblication (i.e. *early mentions*), ordered by year of publication.

Figure 3: Newspaper and blog mentions of retracted articles.

Panel A: News mentions (N=135 retractions) Panel B: Blog mentions (N=365 retractions)

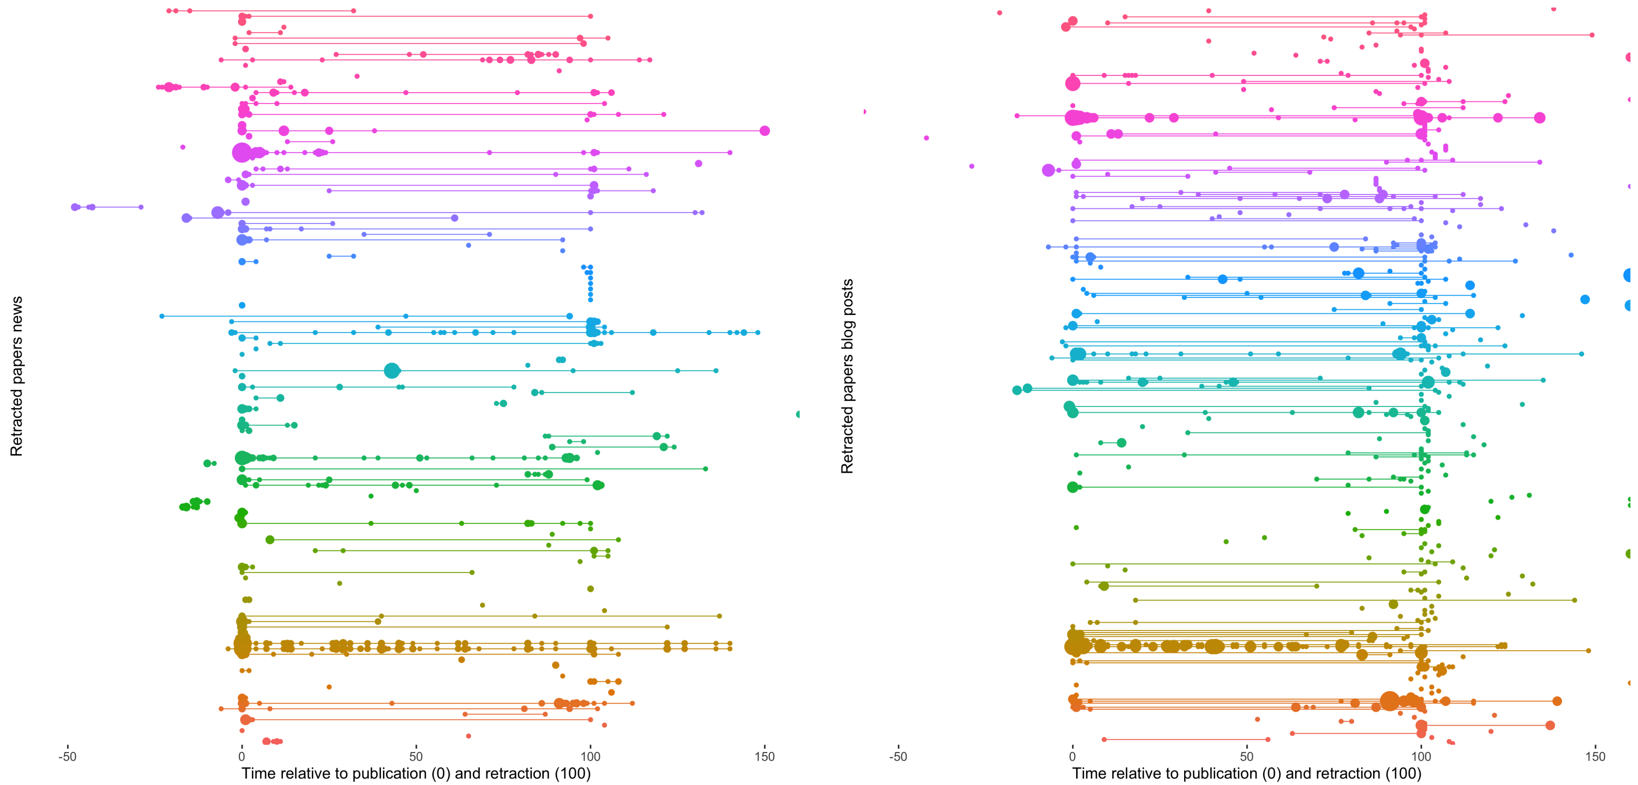

Note: Each line connects the first to the last mention of a single research article on either newspapers (Panel A) or blogs (Panel B) within the considered time window. Dots represent the number of mentions at a certain point in time. The window of analisis focuses on two events: the paper publication date (indexed with 0) and the paper retraction date (indexed with 100). The time score is allocated following the formula $\frac{(t_{mentionposted} - t_{publication})}{(t_{publication} - t_{retraction})} * 100$. The sources of publication date and retraction date are *Altmetric* and *RetractionWatch* respectively.

3 Data and Method

3.1 Conceptual framework

This work studies the possibility that media coverage influences scientists' awareness and assessment of research findings as proxied by citations post-retractions. In other words, whether online attention helps or undermines the process of auto-correcting science.

To understand the interplay between the retraction process of a paper and the information available online one needs to consider how scientific publications feature in the media and what challenges this poses in terms of identification.

A research article that is accepted for publication may endogenously attract media coverage. Online attention may depend on factors such as the salience of a topic, the importance of the findings, the prestige of authors and publishers, the presence of a press release (Sumner et al., 2014, 2016).

Media coverage can bring publicity to a paper which can increase future citations (Phillips et al., 1991; Fanelli, 2013) as well as prompt higher scrutiny from the scientific community making any fault more likely to emerge. Online attention can also inform about the fate of an article as in the case of an expression of concern or a retraction (Serghiou et al., 2021), news that could reach those unaware scientists that would otherwise incorrectly cite a flawed article.

My main intent is therefore to identify the impact of online coverage of an article to its citations post-retraction as a proxy for scientists' opinion over the validity of a piece of research. Holding constant other factors, a loss in citations reflects an erosion of trust in the scientists' work by the scientific community.

To best capture a change in future citations one needs to compare a retracted paper to a similar one which was never retracted and could best approximate the fate of the former in absence of a retraction shock. The treatment and control group similarities not only allow for a more accurate comparison but also account for the potential problem of selection into treatment due to the fact that flaws in prominent papers may be easier to spot.

3.2 Data Overview

This study relies on multiple data sources. First, the treatment sample is extracted from the *RetractionWatch* database[12] which Brainard (2018) defines as "the largest-ever database of retracted articles". The dataset contains a list of retracted research articles[13] together with the following information: title, doi, date of publication, date of retraction, journal, name of authors and their institutions, list of reasons for retraction, and when available, a link to the associated blog post reporting on the paper background story.
I further select papers featuring in either *Scimago* or *Google scholar* rankings. The selected journals appear either as one of the ten highest ranked in Scimago in any of the available subjects or among those listed in Google scholar top publications in any of the existing categories.
Yearly citations are the main outcome of this study and are collected for each article using *Scopus*, one of the two largest abstract and citation database of peer-reviewed literature.[14]
Data on online mentions are gathered thanks to *Altmetric*, a company that since 2011 tracks the attention that research outputs receive online.[15] For each paper I retrieve the *Altscore*, an aggregate measure of online mentions (i.e. it combines all mentions across outlets giving a higher weight to outlets such as newspapers, see appendix Table A.1), and details about single mentions (e.g. date, url, author, title, summary).
Finally, I obtained data on the content of citation statements quoting the research articles in the sample with the support of *Scite.ai*, a recently launched start-up that uses text analysis to categorize reference statements. For each pair of citing and cited study, statements are categorised as "mentioning", "contrasting" and "supportive".[16] In addition, access is gathered for any statement containing the

[12]Version obtained in July 2020.

[13]Dense since the '80s.

[14]For the period considered in the analysis, there exists little difference between Scoups and WoS in terms of coverage (see: https://www.internauka.org/en/blog/scopus-vs-web-of-science). Scopus though has the advantage of having an API easily accessible via *rscopus*, a library by John Muschelli available on R.

[15]I here focus on sources with the highest number of mentions (i.e. newspapers, blogs and Twitter) though Altmetric collects mentions from numerous additional outlets (e.g. Pubpeer, Wikipedia).

[16]According to Rosati (2021)

words "*etract*" or "*ithdraw*".[17]

3.3 Treatment group

The full RetractionWatch database counts $N_r = 21,968$ retractions starting from the 1980 from which only research articles[18] with non missing paper and retraction notice DOI[19] were maintained. To simplify data collection, I then focused on articles published in journals featuring in either Google scholar top publication journals by field or among the ten highest ranked journals in Scimago per subject category.[20] Remaining papers are certainly relevant for the scientific community, hence, it is important to study whether in this case the disinformation is halted or fostered by media coverage. In addition, publications in reputable journals may be more likely to attract media coverage, thus, helping identification. Next, I excluded articles for which I cannot find in Scopus any author with at least one publication in the 9 years before the retraction event.[21] Provided that data availability on research mentions online is only relatively recent, I finally selected retracted papers both published and eventually retracted after 2010, leading to a final sample of $N_r = 990$ retractions for which an appropriate control was found.[22]

3.4 Control group

Trends in citations vary across disciplines, age and media coverage, hence, control publications were selected to mimic pre-retraction characteristics of the treated. This strategy draws from the approach first used in the literature on retractions by Furman et al. (2012) and further developed by Lu et al. (2013) and Jin et al. (2019). The main assumption is that treated papers would continue to perform similarly to control ones in absence of a retraction event.

[17]Manually checked to exclude any false positive.
[18]Excluding for examples: conference abstracts and clinical studies, $N_r = 11,005$.
[19]$N_r = 9,827$.
[20]$N_r = 1,931$.
[21]This is important for an extension of this work in which consequences on authors previous publications are also identified, $N_r = 1,729$.
[22]In 18 cases it was not possible to find a control satisfying the conditions of the control matching algorithm explained in section 3.4.

The selection of the control group proceeds in steps. For each retracted paper I search in Scopus for studies[23] published in the same journal and year of the treated.[24] For each retracted i and potential control pair j I compute the measures listed below.

- *Absolute arithmetic distance* (AD) in citations.

$$AD = \left| \sum_{t=pub}^{retr-1} (c_{it} - c_{jt}) \right|;$$

- *Euclidean distance* (ED) in citations.

$$ED = \left[\sum_{t=pub}^{retr-1} (c_{it} - c_{jt})^2 \right]^{1/2};$$

where c_i indicate the citations paper i receives in year t in the time span between the year of publication pub and the year of retraction $retr$. These measures capture the disparity in citation trends in different ways. AD allows for positive and negative yearly differences to balance over time while any discrepancy is accumulated over time in the case of ED.

- *Early mentions absolute distance* (MD) of blog b and newspaper n mentions whitin two weeks from publication.

$$MD_b = \left|(b_{i,2w} - b_{j,2w})\right| \quad \& \quad MD_n = \left|(n_{i,2w} - n_{j,2w})\right|$$

The reason for choosing a two weeks cutoff draws from observing that notable studies attract most online publicity around their publication as suggested by Figure 3 for treated papers, and it is even more evident for control papers (see Figure A.3). On the other hand, a flawed article, that will eventually be retracted, may later prompt critical mentions. To capture whether a contribution is newsworthy, without including mentions that may relate to the treated article misfortune, I hence focus on a two weeks cutoff from publication date.[25]

[23] Articles or reviews.

[24] N_c = 586.281 overall results.

[25] This threshold is also less sensitive to publication date imprecisions and allows to select controls with a smaller MD as compared to shorter cutoff.

I then retain for each i all j with $AD \leq 10$; $MD_b \leq 10$; and $MD_n \leq 10$.[26] I rank the remaining j in terms of smallest $MD_b + MD_n$ and select two controls (or one depending on availability) with the minimum ED among those. This final selection leads to a sample of $N_c = 1969$ control articles.
The quality of selected controls is assessed in Figure A.1 and A.2 of the appendix. The Euclidean distance between the selected controls and the treated paper is dense around zero (in over 68% of the cases this selection yields a perfect match), and the arithmetic distance is fairly centred around zero. When comparing the two groups distribution of pre-retraction cumulative citations no significance difference emerges, if anything future retracted papers are marginally less likely to have no pre-retraction citations. Inspecting the distribution of early mentions across groups for both newspaper articles and blog posts no significant difference emerges, even though the vast majority of published articles does not appear in either outlets at the time of publication.

3.5 Selected summary statistics

Table 1 illustrates a set of distinct summary statistics for treatment and control group. The top of the table looks at variables which should be similar across the two groups for the identification strategy to be successful. ED and AD are on average somewhat close to zero (0.93 and 0.17 respectively) and both groups of papers attracted an average of about 7 citations in the pre-retraction period, substantially confirming the finding reported in Figures A.1 and A.2. Within two weeks from publication papers experience comparable online mentions on newspapers and blogs, even though eventually retracted papers have on average moderately higher coverage (1.04 vs. 0.79 news articles, and 0.24 vs 0.15 blog posts). The age for the two groups of papers is almost identical by construction. Moving to the bottom of the table one can observe that papers take on average two years to be retracted. Furthermore, yearly citations have a distribution that is very skewed, with 32.2% observations actually equal to 0, a Poisson model would therefore better approximate the distribution of the dependent

[26] These cut offs allow to maximise the number of matches while limiting the maximum conceded distance in either citations or media mentions. These thresholds improve the quality of matches without affecting results.

variable. Unsurprisingly, treated papers cumulate substantially less citations over the years as compared to controls (16.8 vs. 33.9 respectively), but attract generally higher online attention with an *Altscore* of 37.5 for retracted papers and 19.1 for controls. In general, a non negligible share of articles experiences some online coverage, most articles are mentioned on social media (60% of retracted papers and 44% of controls) while only a limited fraction appears in newspaper articles (13% and 12% respectively), in addition blogs actively mention over one third of retracted papers while significantly less attention is devoted to controls. Finally, about one tenth of papers in either group appears in either newspapers or blogs around the publication date.

Table 1: Selected summary statistics

	Mean	Sd	Min	Max	Mean	Sd	Min	Max
	TREATMENT (N=990)				CONTROL papers (N=1969)			
BALANCING VARIABLES								
Euclidean distance					0.937	2.661	0	45.51
Arithmetic distance					0.171	1.699	-10	10
Cum. (no self) citations $(t-1)$	7.103	19.11	0	254	6.807	18.65	0	258
Early news mentions	1.037	7.844	0	134	0.787	5.790	0	127
Early blog mentions	0.240	1.511	0	28	0.152	0.862	0	21
Age	5.138	2.612	0	9	5.140	2.612	0	9
ADDITIONAL VARIABLES								
Time to retract	2.067	2.021	-0.504	9.353				
Yearly citations (no self)	2.628	4.442	0	56	5.259	12.27	0	354.8
Cum. (no self) citations	16.83	30.50	0	418	33.91	75.29	0	1,774
Altscore	37.54	274.5	0	7,128	19.09	130.2	0	3,728
Tweeters count	32.12	349.9	0	10,105	13.55	165.4	0	5,100
News count	1.459	8.589	0	122	1.068	6.158	0	113
Blog count	0.871	3.283	0	65	0.287	1.333	0	27
Any social media mention	0.597	0.491	0	1	0.443	0.497	0	1
Any news mention	0.136	0.343	0	1	0.119	0.324	0	1
Any blog mention	0.369	0.483	0	1	0.110	0.313	0	1
Any early visibility	0.111	0.314	0	1	0.0945	0.293	0	1

Note: Self-citations are excluded from citation count. *Early mentions* include all news and/or blog posts published within 2 weeks from publication. *Altscore* is a weighted average of all online mentions across outlets. Media *counts* are the number of outlets/accounts referring to a paper at any point in time. All papers are published/retracted between 2011 and 2020.

3.6 Empirical Approach

The study employs a difference-in-differences strategy that allows to compare the evolution of citations of retracted papers before and after retraction relative to citations of a control group of non-retracted studies published in the same journal and year and with a comparable trend in yearly citations before retraction. Treatment and control papers also have similar number of online mentions (on blogs and newspapers) within two weeks from the day of publication (i.e. *early visibility*) to account for unobservable characteristics which make a study newsworthy and could therefore create a problem of selection into retraction. Therefore, the regression model is the following:

$$E(Y_{igt}) = exp(\alpha + \beta_1 Post_{gt} + \beta_2 T_i * Post_{gt} + \beta_3 T_i * Post_{gt} * Media_i + \\ + \beta_4 Post_{gt} * Media_i + \delta_i + f(age_{it}) + \delta_\tau) \quad (1)$$

where i is the treatment (or control) paper, g is the case-level group and includes the retracted paper and its respective controls, t are years relative to the retraction. The dependent variable Y represent a paper yearly citation count and exclude self-citations, as the estimation wants to capture the reaction of the scientific community other than that of the authors involved. *Post* is an indicator variable equal to one for all years since retraction, T is an indicator for retracted articles, and *Media* is an indicator capturing whether an article was exposed to online coverage. Different media dummies will be used to indicate articles with at least one online mention within two weeks from publication in newspapers and/or blogs (i.e. early visibility) or at least one overall mention in any of the media outlet analysed (i.e. any socialmedia, newspaper articles or blog posts). In order to look at different level of media exposure of each paper, indicators are also derived from the distribution of *Altscore*, an aggregate measure of weighted online mentions. Fixed effects are included for each paper δ_i and each calendar year δ_τ while $f(age_{it})$ represents a full set of dummies for years since publication (age) and is meant to flexibly control for the age of the articles.[27] The coefficient β_2 captures the effect of a retraction shock on citations of retracted papers as compared to similar control papers. The coefficient β_3

[27]Note that the interaction term $T_i * Media_i$ is absorbed by the paper fixed effect.

captures any difference in the effect of the shock for papers that received media attention. To look at the dynamics of the differential effect of *Media*, estimates will be presented for a model that replace the indicator *Post* with a full set of dummies for each year relative to the year of retraction.[28] Given the skewed nature of the dependent variable, I follow a long-standing tradition in bibliometric studies, hence I use a pseudo Poisson regression model developed by Correia et al. (2020)[29] where consistency is achieved under the only assumption that the conditional mean of the dependent variable is correctly specified (Gourieroux et al., 1984). Finally, standard errors are clustered at the case g level.

4 Results

Table A.2 shows results for a simple difference-in-differences analysis for the pooled sample of retracted papers and selected controls. Estimates imply that relative to controls, retracted papers experience a 65%[30] loss in yearly citations after the shock and the magnitude is comparable to previous studies (Furman et al., 2012; and Azoulay et al., 2015) which rely on different samples, disciplines and time periods.[31] Figure A.4 illustrates the dynamic of the effect of a retraction. The post-retraction loss in citations increases over time and there is no evidence of pre-trends.[32]

4.1 Main results

Table A.3 to A.5 report results from the main specification. The tables differ by measures of media coverage, using indicators for papers with at least one mention within two weeks from publication (early visibility), papers with at least one mention overall in a certain online outlet (any news, blog or social media)

[28] $E(Y_{ig\tau}) = exp(\sum_{t=r-4}^{r-2} \beta_{1t} * d_t + \sum_{t=r}^{r+6} \beta_{1t} * d_t + \sum_{t=r-4}^{r-2} \beta_{2t} * d_t * T_i + \sum_{t=r}^{r+6} \beta_{2t} * d_t * T_i + \sum_{t=r-4}^{r-2} \beta_{3t} * d_t * T_i * Media_i + \sum_{t=r}^{r+6} \beta_{3t} * d_t * T_i * Media_i + \sum_{t=r-4}^{r-2} \beta_{4t} * d_t * Media_i + \sum_{t=r}^{r+6} \beta_{4t} * d_t * Media_i + \delta_i + f(age_{i\tau}) + \delta_t)$

[29] http://scorreia.com/software/ppmlhdfe/

[30] In order to interpret the estimated coefficient as percentage loss in citations I use the following transformation: $1 - exp(-1.06) = 0.65$.

[31] Estimates are similar when using an IHS (Inverse hyperbolic sine) transformation of the dependent variable.

[32] Note that effects in the year of retraction are also small due to the fact that papers in the sample get retracted at different points within the year.

or papers that fall in some part of the Altmetric score (Altscore) distribution. The tables highlight the difference-in-differences coefficient $Post * Treatment$, according to which the average citation penalty of a paper after its retraction amount to 56-62% across all specifications. The relative effect for papers that experienced some media coverage is estimated by the triple interaction $Post * Treatment * Media$.[33] Retracted papers which attracted media coverage experience an additional penalty in post-citations of about 28-36% and the effect seems monotonically increasing in the amount of coverage received (see Table A.5). The almost entirety of these estimates is highly significant. Figure 4 represents the dynamics of the additional penalty in presence of (alternative measures of) media coverage. The loss in yearly citations becomes progressively more evident over time without any sign of recovery, and I find no evidence of pre-trends.

4.2 Robustness checks

4.2.1 Including retraction year into *Post* indicator

Previous estimates illustrate effects on citations for all years strictly after the one of retraction (i.e. excluding the year of retraction). The rationale behind this choice is the fact that papers can get retracted at any point during the year and this can therefore act as a confounder.[34] Nonetheless, Tables A.6 to A.8 show that the main results are not sensitive to this decision. If anything, the additional effect of early visibility is smaller in case of blog mentions (see Table A.6 column (2)). This difference may be interpreted as a sanity check, given that most blog mentions appear later when the paper gets retracted. In addition, the fact that early visibility effects are less significant, speaks to a possible information effect of the media coverage which emerges more clearly at a later stage and it is better captured by overall measures of online coverge as in A.7 and A.8.

[33]Where the *Media* variable is defined in alternative ways across specifications as described at the top of this paragraph.

[34]Figure A.4 and Figure 4 show smaller or insignifican effects in the year of retraction relative to the previous year.

Figure 4: Dynamics of retracted papers penalty with media coverage

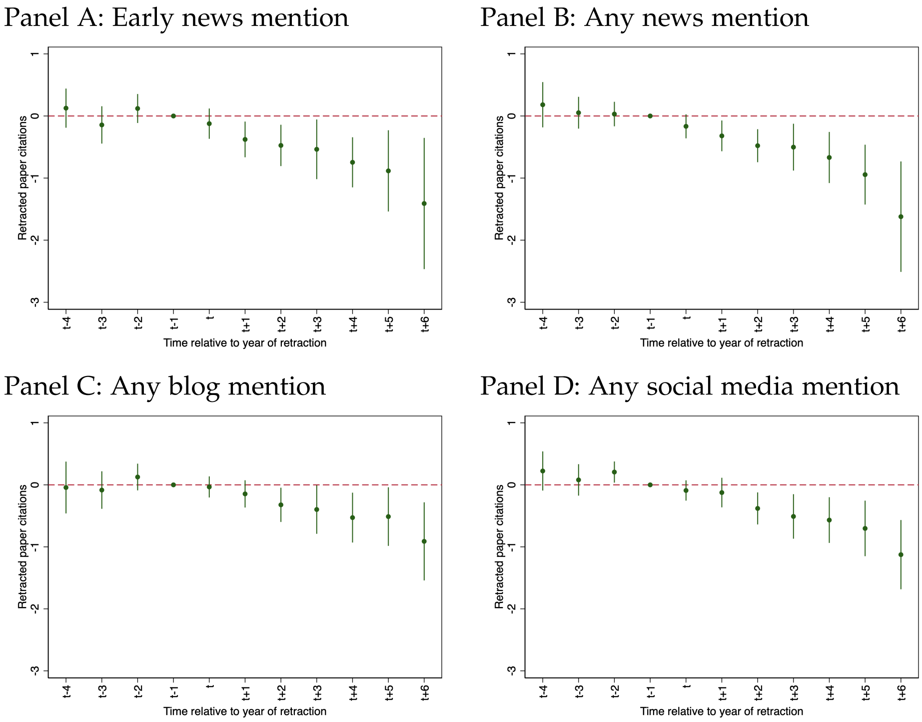

Note: Estimates replicate the following models: Table A.3 column (3) for Panel A; Table A.4 column (3)-(4)-(1) respectively for Panel B, Panel C and Panel D. Models are estimated replacing the *Post* indicator with a full set of dummies for each year relative to the retraction ($t-1$ excluded). The coefficients displayed are that of the interaction between time dummies, a treatment indicator and a media indicator while vertical lines represent 95% CI.

4.2.2 Actively cited papers

The algorithm for selecting controls attempts to choose papers that could likely mimic the citation path of retracted papers absent the retraction shock. Finding good controls for retracted papers that are not actively cited soon after publication may be challenging and could bias estimates. For this reason I here exclude all retracted papers (and respective controls) with zero citations in any year before retraction. This exercise halves the original sample.[35] Even so, Tables A.9 to A.11 confirm the results all remain robust.

[35] 48% observations left in either treated or control group.

4.3 Selection into retraction

In section 3.1 I argued that a challenge one faces when trying to understand the interaction between the retraction process and media coverage, arises from the endogeneity of the latter. To circumvent this issue to some extent, I turn to the text analysis of titles of research articles. This in turn allows me to use the presence of specific words to control for papers' endogenous coverage.

More specifically, I start with the full sample of eventually retracted articles published (and retracted) after 2010 and for each of these articles I add to the sample twenty randomly selected articles that appear in the same journal and year but were never retracted.[36] I then use the titles of these papers as corpus of analysis.[37] After cleaning the text according to Porter (1980) algorithm, Figure 5 shows the most frequent words present in the titles of papers that experience some (Panel A) or no (Panel B) online coverage (in newspapers or blogs) within two weeks from publication. On the one hand, popular papers mention more often words shuch as "cancer", "patient" and "disease", on the other, articles that did not feature in the media often quote different words such as "model" or "system". In what follow I try using this differences to predict articles coverage.

After building the document-term matrix of words (unigrams and bigrams) that appear in at least 100 titles I randomly split the observations into 90% training and 10% testing subsample. The training sample is used to select words with some predictive power for papers' media coverage based on lasso selection procedure. The testing sample is then used to compute the out-of-sample performance of the predicted media coverage based on the selection.[38]

The lasso estimates and the set of selected variables (words) depends on the penalty level λ. I obtained alternative lists of selected words using different pro-

[36] This selection facilitate a speedy computation without restricting the corpus of titles. Among the 1008 retracted papers in the sample, 44 have less than 20 associated random controls due to the respective scarsity of potential controls found in Scopus.

[37] $N = 20755$

[38] The lasso estimation minimizes the mean squared error subject to a penalty on the absolute size of coefficient estimates and where λ controls the overall penalty level.

$$\hat{\beta}_{lasso} = argmin \frac{1}{n} \sum_{i=1}^{n} (Media_i - \sum_{j=1}^{p} \beta_j Word_{ij})^2 + \frac{\lambda}{n} \sum_{j=1}^{p} \psi_j \mid \beta_j \mid$$

Due to the nature of the penalty, the lasso sets some coefficients exactly to zero and in doing so removers some predictors from the model.

Figure 5: Papers' titles wordclouds

Panel A: Title of papers w media (N=1961) Panel B: Title of papers w/o media (N= 18794)

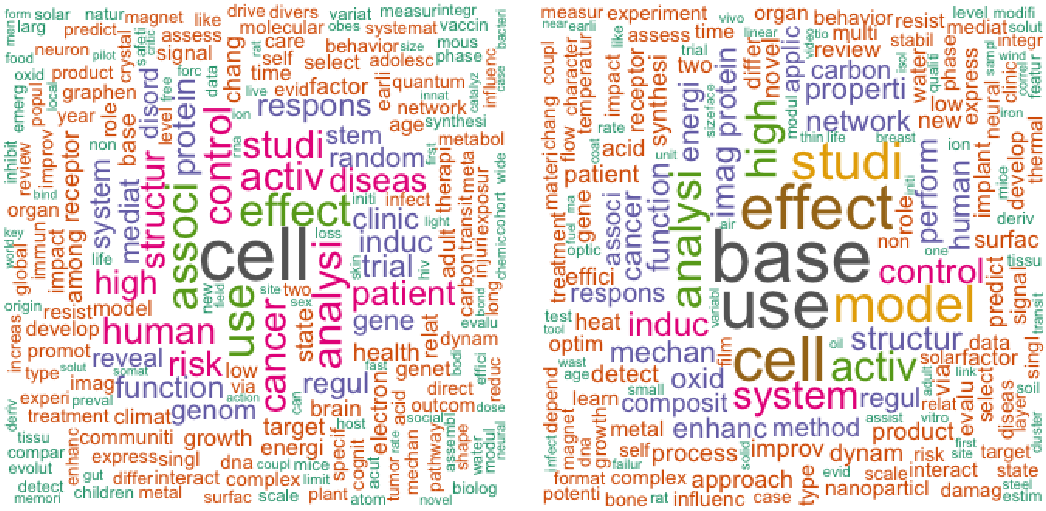

cedures that choose the optimal penalty level using: (a) EBIC information criteria; (b) AICC information criteria; (c) K-fold cross-validation and (d) Rigourous (theory-driven) penalty levels. These procedures are then repeated including a full set of subject fixed effects, publication year fixed effects and excluding retracted articles from the sample. This strategy allows me to then estimate the following model:

$$Retraction_{ijp} = \beta_1 Media_{ijp} + \beta_2 \widehat{Media}_{ijp} + \delta_j + \delta_p + \epsilon_{ijp} \tag{2}$$

where for each article i published in year p and journal j, *Retraction* is an indicator for whether the article was retracted, *Media* is a dummy taking value one if the article gained any online coverage (in either newspapers or blogs) within the first two weeks from publication, while δ_j and δ_p absorb journal fixed effects and publication year fixed effects respectively. Estimating the *Media* impact on the likelihood of a retraction (β_1) is challenging as it is difficult to exclude that researchers may choose to investigate salients topics that, given their relevance, are scrutinized differently from the scientific community (see for example Serra-Garcia and Gneezy, 2021) leading to different retractions rates, despite the fact that these topics may be of interest to the general public and hence attract media coverage.

Table 2: Selected words and media coverage

	Media coverage							
	Linear				Logit			
	EBIC	AICC	CV	Rigourous	EBIC	AICC	CV	Rigourous
# of adult	0.110***	0.093***	0.092***		0.067***	0.044**	0.045**	0.073***
	(0.036)	(0.035)	(0.035)		(0.018)	(0.017)	(0.017)	(0.017)
# of algorithm		-0.047***	-0.048***	-0.049***		-0.186**	-0.185**	
		(0.008)	(0.008)	(0.007)		(0.076)	(0.076)	
# of brain	0.081***	0.068**	0.066**		0.062***	0.041**	0.041**	
	(0.026)	(0.027)	(0.027)		(0.015)	(0.017)	(0.017)	
# of climat	0.228***	0.221***	0.219***		0.134***	0.120***	0.126***	0.132***
	(0.054)	(0.050)	(0.050)		(0.022)	(0.020)	(0.020)	(0.021)
# of commun	0.103***	0.091***	0.089***		0.073***	0.061***	0.061***	0.074***
	(0.034)	(0.034)	(0.035)		(0.019)	(0.020)	(0.020)	(0.019)
# of composit	-0.046***	-0.036***	-0.035***	-0.051***	-0.124***	-0.098***	-0.099***	-0.127***
	(0.006)	(0.007)	(0.007)	(0.007)	(0.034)	(0.033)	(0.033)	(0.034)
# of disord	0.170***	0.162***	0.162***		0.089***	0.080***	0.081***	0.094***
	(0.058)	(0.057)	(0.057)		(0.022)	(0.021)	(0.021)	(0.023)
# of earli	0.090***	0.078**	0.077**		0.066***	0.057***	0.056***	
	(0.034)	(0.033)	(0.033)		(0.019)	(0.019)	(0.019)	
# of genom	0.076**	0.065**	0.065**		0.045**	0.037**	0.036**	0.044**
	(0.033)	(0.033)	(0.033)		(0.017)	(0.017)	(0.017)	(0.017)
# of global		0.070**	0.071**			0.049***	0.049***	
		(0.033)	(0.033)			(0.019)	(0.019)	
# of graphen		0.076**	0.075**			0.087***	0.087***	
		(0.031)	(0.031)			(0.022)	(0.022)	
# of meta_analysi	0.098**	0.120***	0.120***		0.061***	0.080***	0.061***	
	(0.039)	(0.040)	(0.042)		(0.019)	(0.027)	(0.022)	
# of model	-0.043***	-0.036***	-0.037***	-0.044***	-0.074***	-0.065***	-0.064***	-0.075***
	(0.006)	(0.007)	(0.007)	(0.007)	(0.016)	(0.016)	(0.016)	(0.016)
# of network_ETX		-0.063***	-0.063***	-0.061***	-0.183**	-0.181**	-0.181**	
		(0.011)	(0.011)	(0.010)	(0.081)	(0.080)	(0.080)	
# of neuron	0.082**	0.070*	0.069*		0.056***	0.041*	0.039*	
	(0.038)	(0.038)	(0.038)		(0.020)	(0.021)	(0.021)	
# of reveal	0.110***	0.102***	0.102***		0.070***	0.065***	0.059***	0.070***
	(0.029)	(0.029)	(0.029)		(0.016)	(0.015)	(0.015)	(0.016)
# of risk	0.080***	0.068**	0.067**		0.054***	0.034**	0.034**	0.056***
	(0.028)	(0.027)	(0.028)		(0.015)	(0.015)	(0.015)	(0.015)
# of stem	0.092***	0.081***	0.084***		0.060***	0.051***	0.050***	0.059***
	(0.030)	(0.030)	(0.030)		(0.019)	(0.019)	(0.019)	(0.019)
# of STX_structur	0.118***	0.113***	0.112***		0.079***	0.081***	0.079***	
	(0.035)	(0.035)	(0.034)		(0.019)	(0.018)	(0.018)	
# of trial	0.105**	0.074	0.075		0.074***	0.048	0.048	0.078***
	(0.050)	(0.054)	(0.054)		(0.027)	(0.032)	(0.032)	(0.027)
# of vitro		-0.052***	-0.054***	-0.069***		-0.182***	-0.180***	
		(0.009)	(0.009)	(0.008)		(0.058)	(0.058)	
Total # of words	-0.006***	-0.005***	-0.004***	-0.004***	-0.006***	-0.004***	-0.004***	-0.006***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Citations $year_p$	0.018***	0.017***	0.017***	0.018***	0.014***	0.013***	0.013***	0.014***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.001)	(0.001)	(0.001)	(0.001)
N	20755	20755	20755	20755	20755	20755	20755	20755
Out-of-sample R2	0.095	0.101	0.100	0.077				
R-squared	0.088	0.114	0.115	0.074				
Out-of-sample accuracy	63.47	61.88	60.10	71.08	66.84	63.76	67.37	73.49
Overall accuracy	70.50	63.92	62.25	65.26	76.37	69.36	62.53	71.68

Note: Estimates from OLS (columns 1-4) or Logit regression (columns 5-8). The dependent variable *Media coverage* is an indicator for whether a paper attracted online coverage at publication. In column (1)-(4) predictors are selected based on *lasso* while column (5)-(8) predictors are selected based on *lassologit*. Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

The inclusion of $\widehat{Media} = \sum_s \hat{\beta}_{s,lasso} SelectedWord_s$ as predicted from the lasso procedure, where *SelectedWord* represents the number of times a selected n-gram appears in the title of a paper i, allows me to control for endogenous topic selection that could otherwise lead to bias. [39]

Table 2 shows the correlation between some of the most powerful lasso selected predictors and the *Media* indicator variable. The n-grams with the largest coefficients provide insights into which articles receive media coverage. For example, the word "climate" appears. Similarly, the n-grams "brain", "graphen", "genom" and "stem" all represent research topics of large interest. Also, some research methodologies seem popular as suggested from the n-gram "meta analysis" and "trial". Accuracy ranges between 60 and 76% across procedures and more parsimonious lasso (and logit lasso) seem to provide better performing selections. The fraction of correctly classified observations reaches up to 86% when a full set of subject and year fixed effects are included and when retracted papers are excluded.[40]

Table 3: Likelihood of retraction and media coverage

	Panel A: Retraction									
	OLS		EBIC		AICC		CV		Rigourous	
Media coverage	0.009**	0.019***	0.014***	0.021***	0.016***	0.021***	0.016***	0.022***	0.013***	0.020***
	(0.004)	(0.007)	(0.005)	(0.007)	(0.005)	(0.007)	(0.005)	(0.007)	(0.005)	(0.007)
Predicted media			-0.078***	-0.074***	-0.067***	-0.068***	-0.066***	-0.068***	-0.135***	-0.140***
			(0.025)	(0.028)	(0.017)	(0.022)	(0.017)	(0.021)	(0.043)	(0.054)
	Panel B: Retraction									
			EBIC		AICC		CV		Rigourous	
Predicted media			-0.059***	-0.065**	-0.048***	-0.060***	-0.048***	-0.068***	-0.104***	-0.127**
			(0.022)	(0.028)	(0.015)	(0.021)	(0.015)	(0.021)	(0.039)	(0.054)
Publication year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Journal FE	N	Y	N	Y	N	Y	N	Y	N	Y
N	20755	20755	20755	20755	20755	20755	20755	20755	20755	20755
N clusters	1008	1008	1008	1008	1008	1008	1008	1008	1008	1008

Note: Estimates from OLS regression. The dependent variable *Retraction* is an indicator for whether a paper was retracted. *Media coverage* is an indicator for whether a paper attracted online coverage at publication. *Predicted media* is media coverage as predicted from the respective *lasso* procedures. Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

[39] Summary statistics of main variables and a selection of n-grams are displayed in Table A.12.

[40] The most powerfull predictors selected with these alternative strategies remain fairly similar (not shown and available upon request).

Equation (2) estimates are reported in Table 3 (Panel A) where despite the differences in n-gram selection and predictive accuracy across models, very similar results emerge across all specifications (mirroring results are displayed in Table A.13 for logit estimations and in Table A.14 where lasso procedures are trained within subjects and years and excluding retracted articles). The evidence suggests that articles with higher *predicted* media coverage are less likely to experience a retraction. The interpretation of this result is twofold. On the one hand, the fact that popular articles are retracted less often may seem reassuring, on the other, it may indicate that more "interesting" research articles may be reviewed with a laxer standard (as suggested in Serra-Garcia and Gneezy, 2021). Under the assumption that predicted media coverage effectively controls for endogenous topic selection, the remaining variation in media coverage is arguably exogenous and therefore allows to estimate the impact of additional attention. Hence, estimates show that wider media coverage at publication, conditional on its prediction, leads to higher chances of retractions, but the magnitude of this effect remain small. This result justifies selecting controls with early media presence similar to that of their retracted counterpart as allowing the however small selection into treatment of more popular articles would otherwise bias the main results reported in section 4.1. Finally, one could be concerned about the common inclusion of both the media indicator and its text-based prediction due to their positive correlation ($\rho \approx 0.3$). To this respect, Table 3 additionally reports the impact of the two regressors separately (see Panel A column (1-2) and Panel B respectively), the magnitudes of coefficients varies only slightly in this case, ressuring us against a collinearity issue.

4.4 Journal visibility

In the following section I offer an additional way to circumvent the endogeneity of media coverage. In what follows I argue that non-retracted articles published in the same journal and year as a retracted one may attract online coverage which is arguably *exogenous* to the retracted article own visibility. Based on this, a good proxy for online visibility of a specific journal and year is the average visibility of all non-retracted papers published in there.[41]

$$Visibility_{jp} = \frac{1}{n} \sum_{k \neq i} Altscore_{kjp} \tag{4}$$

where k are non-retracted papers published in same journal j year p as the retracted paper i. Alternatively I use the average share of $k \neq i$ published in j and p with some media mentions. Hence, I can study the following relationship using an OLS regression in a cross-sectional context:

$$Y_{ijp} = \beta Visibility_{jp} + \delta_p + \nu X_{ijp} + \epsilon_{ijp} \tag{5}$$

where for each retracted paper i published in j in year p, Y represents either one of the dependent variables: $Time\ to\ retract = (Retraction\ date - Publication\ date) \times \frac{12}{365}$ or $DiD = [E(cit_1^T) - E(cit_1^C)] - [E(cit_0^T) - E(cit_0^C)]$ the individual loss in citations obtained comparing each retracted paper to its selected controls for different pre- and post- time windows. In addition, δ_p indicate publication year fixed effects, while X_{ijp} control for $N_{jpk \neq i}$ number of non-retracted papers in same journal-year as retrieved from Scopus, $\frac{1}{n} \sum_{k \neq i} ED_{kjp}$ their average Euclidean distance in citation from the retracted one, and $\sum_{p \leq t < r} cit_{ijpt}$ cumulative citation of i before retraction year r. Standard errors are clustered at the journal level.

4.4.1 Time to retract

Figure 6 shows that papers are retracted faster when published in journals where the average article attracts higher online coverage. Table A.15 (column (1)) illustrates that one standard deviation increase in journal visibility (measured as the average Altscore of non-retracted articles in a journal-year) reduces time to retraction by approximately 15% of its average. Looking across the remaining columns, the relationship is robust to different measures of visibility.

[41] The measure is based on the entire pool of papers published in same year and journal as the retracted ones which where extracted from Scopus as a first step in the selection of potential controls for the main analysis.

Figure 6: Months to retraction and Journal-year average visibility

Panel A: Raw data Panel B: Absorb Controls + FE

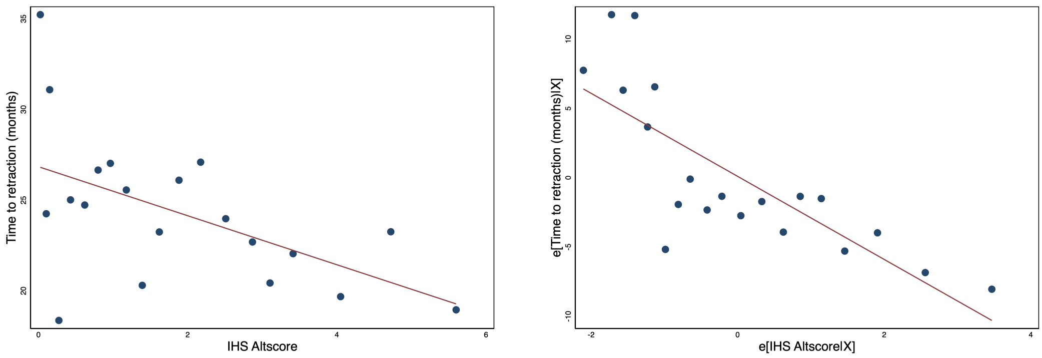

Note: The vertical axis represents the time intercurring between an article publication and its retraction, expressed in months. The orixontal axis represents the inverse hyperbolic sine transformation of journal visibility, measured as the average Altscore of non-retracted papers that appear in the same yournal and year of the retracted one. Controls include the number of non-retracted articles within same journal and year of the treated, the average Euclidean distance of those from the retracted paper, and the level of (non-self) cumulative citations of the retracted paper before retraction. Publication year fixed effects are included. * p<0.1, ** p<0.05, *** p<0.01.

4.4.2 Loss in citation

Figure 7 shows that a retracted paper experiences a significantly larger loss in citations if published in a journal whose average (non-retracted) article receives higher online coverage. Comparing each retracted paper to its selected controls for different pre- and post- time windows, one notice this negative relationship becomes stronger when looking at wider time windows around the year of retraction. These same conclusions are evident in Table A.16 were alternative visibility measures are used. These findings reassure us on the validity of the main identification strategy as they confirm the main results presented in section 4.1.

4.5 Heterogeneity by discipline

Various disciplines have distinctive publication practices which could create different incentives at publication and therefore lead to heterogeneous effect. Table A.17 and A.18 together with Figure 8 illustrate that this may indeed be the case. What consistently emerges across specifications is that, in the case of social

Figure 7: Loss in citation and Controls average mentions

Panel A: Raw data

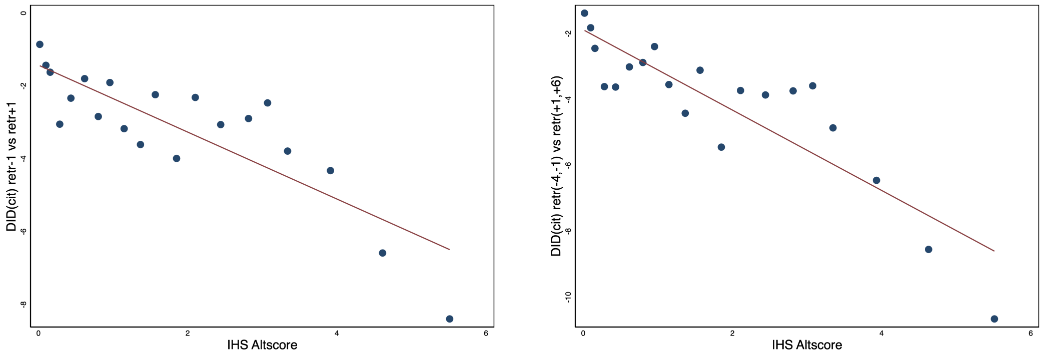

Panel B: Absorb Controls + FE

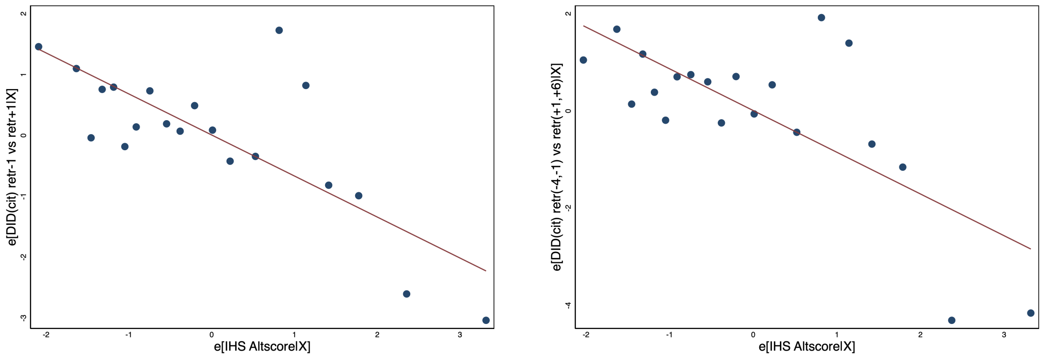

Note: The vertical axis represents the individual loss in citations obtained comparing each retracted paper to its selected controls for different pre- and post- time windows. The time window around retraction become larger moving left to right. The orixontal axis represents the inverse hyperbolic sine transformation of journal visibility, measured as the average Altscore of non-retracted papers that appear in the same yournal and year of the retracted one. Controls include the number of non-retracted articles within same journal and year of the treated, the average Euclidean distance of those from the retracted paper, and the level of (non-self) cumulative citations of the retracted paper before retraction. Publication year fixed effects are included.

sciences,[42] there is no additional penalty associated to retracted papers with media attention. Perhaps one interpretation is that the timing of publications in hard sciences is generally fast; while working papers in social sciences may circulate for longer inside the scientific community. In the latter case, media coverage may therefore offer little room for update on the validity of the study as compared to the former case where online attention may further stimulate

[42]Disciplines are identified using Scopus journal classification. Social sciences are: business and technology, humanities and other; while hard sciences are: life sciences, environment, health and physical sciences.

the academic discussion around a paper.[43] One caveat is that the subsamples of disciplines are quite small in particular in the case of social sciences.[44]

Figure 8: Dynamics of retracted papers penalty with media coverage by discipline

Panel A: Hard sciences with visibility

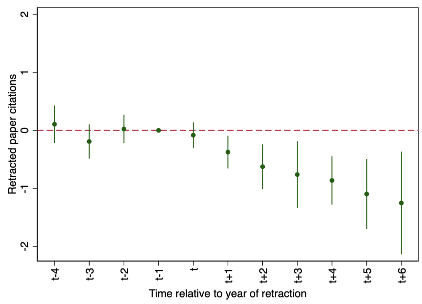

Panel B: Social sciences with visibility

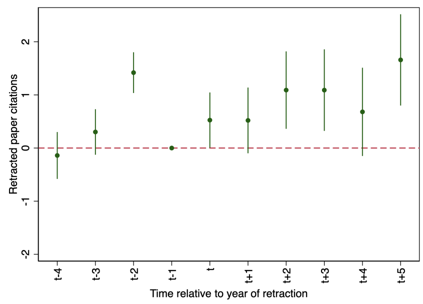

Panel C: Hard sciences with Altscore >p50

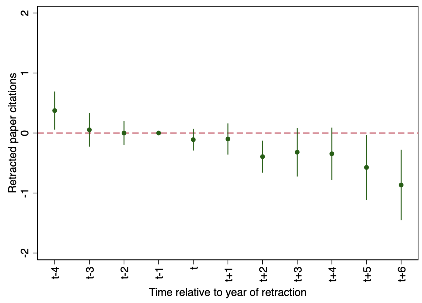

Panel D: Social sciences with Altscore >p50

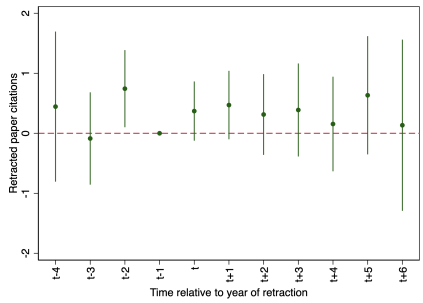

Note: Hard sciences: life sciences, environment, health and physical sciences. Social sciences: business and technology, humanities, other social sciences. Estimates replicate the following models: Table A.17 column (3)-(4) for Panel A and Panel B; Table A.18 column (3)-(4) respectively for Panel C and Panel D. Models are estimated replacing the *Post* indicator with a full set of dummies for each year relative to the retraction ($t-1$ excluded). The coefficients displayed are that of the interaction between time dummies, a treatment indicator and a media indicator, for different subsamples of discipline, while vertical lines represent 95% CI.

[43] Related to this, Wohlrabe and Bürgi (2021) suggests that in the case of economics, the practice of releasing working papers before their publication in a journal has a positive impact on citations.

[44] Over 80% of retraction appears in hard sciences (809) of which 12% (95) with early visibility and 59% (475) with Altscore above median. Of the 179 retraction in social sciences 8% (15) have early visibility and 53% (95) have Altscore above median.

4.6 Citation textual content

One additional exercise is that of looking directly at the textual content of citations. *Scite.ai* (a newly launched platform featuring in Nature)[45] scans article PDFs for references to papers and categorises these references as mentioning, contrasting or supporting.[46] With the platform support, I built a dataset of yearly citation statements for each classification, paper and year and performed an exercise equivalent to that of Section 4.1. Tables A.19 to A.21 substantially corroborate the main findings. Retracted papers experience a penalty in all type of citation statements after the retraction shock, and for citation statements that only mention the study, this penalty is aggravated in presence of media coverage. No additional change is detected for either contrasting or supporting references. One caveat is that almost the entirety of the citation statements is classified as mentioning.

4.6.1 Information mechanism

This works has so far shown that retractions disappear from the literature at a faster pace in presence of media coverage. This additional effect of media may be derived by two different mechanisms: (a) higher scrutiny by the scientific community to a paper that gained publicity; (b) additional information provided to some part of the scientific community which would have otherwise remained unaware of the retraction. Although difficult to distinguish, one way to corroborate the information mechanism is to check whether the content of remaining ex-post citations is more "accurate" in presence of media coverage. With the help of *Scite.ai*, I collected for each retracted paper all yearly citation statements that mentioned the retraction. Citation statements were searched for the terms "*etract*" or "*ithdraw*", manually excluding false positives. I then estimate the following regression model:

$$E(Y_{it}) = exp(\alpha + \beta_1 Post_{it} * Media_i + \delta_i + \delta_t + f(age_{it}) + \delta_\tau) \quad (6)$$

where for each retracted paper i and year relative to retraction t, Y represents

[45] https://www.nature.com/articles/d41586-020-01324-6

[46] The classification is according to Rosati (2021)

the number of citation statements mentioning the paper is retracted, *Post* is an indicator for year strictly after retraction, *Media* is an indicator for whether a paper gained some kind of online coverage. Estimates of β_1 capture the differential change in number of citations "correctly" mentioning the retraction (after the shock) in presence of media coverage. Fixed effects are included for each paper δ_i, each year relative to retraction δ_t, each year since publication $f(age_{it})$ and each calendar year δ_τ. Standard errors are clustered at the retraction level.

Table 4: Citation statements mentioning paper is retracted

	(1)	(2)	(3)
VARIABLES		Statements mentioning paper is retracted	
Post * Early visibility	2.077***		
	(0.457)		
Post * Altscore >p50		1.562**	
		(0.670)	
Post * Altscore 3rd quintile			-0.101
			(1.344)
Post * Altscore 4th quintile			1.137
			(1.223)
Post * Altscore 5th quintile			2.305**
			(1.155)
Article FE	Y	Y	Y
Age FE	Y	Y	Y
Year FE	Y	Y	Y
Relative yr FE	Y	Y	Y
Pseudo R2	0.361	0.341	0.355
N	531	531	531
N clusters	95	95	95
N full	5591	5591	5591

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations statements received by each paper in a particular year which explicitly mention the retraction. Early visibility is an indicator for papers with at least one mention (in newspapers and/or blogs) within two weeks from publication. Altscore is an aggregate measure of weighted online mentions. All models incorporate article fixed effects, a full suite of calendar-year effects, article age indicator variables and dummies for each year relative to the retraction. Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table 4 confirms that the number of references correctly mentioning the cited paper is retracted increases significantly in presence of media coverage. This result support the hypothesis that media coverage provides additional *information* on retractions, hence favouring the belief update of part of the scientific community which would have otherwise remained unaware. One caveat to consider is the small sample of retractions for which an "accurate" yearly-citation is indeed observed (slightly less than 10% of the treated sample).[47]

5 Conclusion

Bad science can be hard to eradicate. Recent studies document that scientific retractions lose a significant amount of citations after being retracted, but worringly still get cited long after they were removed from the literature. This creates the potential for dissemination of misinformation within and outside academia. This work shows that media coverage can positively influence the auto-correcting process of science. I use a conditional difference-in-differences strategy to show that retracted articles experience larger citation-losses in presence of media coverage and the remaining post-citations are more often correct in mentioning the paper is indeed retracted. I further show suggestive evidence of small selection into treatment for papers attracting excess coverage and that journals that generally publish more popular articles are those where retractions appear faster. Finally, the differential effect of media coverage is observed only for hard sciences, suggesting distinct publication practices or different topic salience may impact the visibility of a retraction. While media seems to help the scientific community to update beliefs about the quality of a study, one should think whether this could generate unintended consequences for the main audience of mainstream media: the general public. I indeed show that newspapers, as opposed to blogs, are more likely to publicise the publication of a paper rather than informing about its later retraction. This possible misinformation can impact public perception of scholars' trustworthiness and therefore deserves to be explored with further research.

References

H. Allcott and M. Gentzkow. Social media and fake news in the 2016 election. *J. Econ. Perspect.*, 31(2):211–236, 2017. ISSN 08953309. doi: 10.1257/jep.31.2.211.

[47] This is consistent with previous work by Schneider et al. (2020) which finds that, for the case considered, the retraction is not mentioned in 96% of direct post-retraction citations.

P. Azoulay, J. L. Furman, Krieger, and F. Murray. Retractions. *Rev. Econ. Stat.*, 97(5):1118–1136, dec 2015. ISSN 0034-6535. doi: 10.1162/REST_a_00469. URL `https://direct.mit.edu/rest/article/97/5/1118-1136/58260`.

P. Azoulay, A. Bonatti, and J. L. Krieger. The career effects of scandal: Evidence from scientific retractions. *Res. Policy*, 46(9):1552–1569, 2017. ISSN 00487333. doi: 10.1016/j.respol.2017.07.003.

J. Brainard. Rethinking retractions. *Science (80-.).*, 362(6413):390–393, oct 2018. ISSN 0036-8075. doi: 10.1126/science.362.6413.390. URL `https://www.science.org/doi/10.1126/science.362.6413.390`.

D. Card and S. Dellavigna. What do editors maximize? Evidence from four economics journals. *Rev. Econ. Stat.*, 102(1):195–217, 2020. ISSN 15309142. doi: 10.1162/rest_a_00839.

D. Card, S. Dellavigna, P. Funk, and N. Iriberri. Are Referees and Editors in Economics Gender Neutral? *Q. J. Econ.*, 135(1):269–327, 2020. ISSN 15314650. doi: 10.1093/qje/qjz035.

S. Correia, P. Guimarães, and T. Z. Zylkin. Fast Poisson estimation with high-dimensional fixed effects. *Stata J.*, 20(1):95–115, 2020. ISSN 15368734. doi: 10.1177/1536867X20909691.

E. Dumas-Mallet, A. Garenne, T. Boraud, and F. Gonon. Does newspapers coverage influence the citations count of scientific publications? An analysis of biomedical studies. *Scientometrics*, 123(1):413–427, 2020. ISSN 15882861. doi: 10.1007/s11192-020-03380-1. URL `https://doi.org/10.1007/s11192-020-03380-1`.

D. Fanelli. Any publicity is better than none: newspaper coverage increases citations, in the UK more than in Italy. *Scientometrics*, 95(3):1167–1177, jun 2013. ISSN 0138-9130. doi: 10.1007/s11192-012-0925-0. URL `http://link.springer.com/10.1007/s11192-012-0925-0`.

J. L. Furman, K. Jensen, and F. Murray. Governing knowledge in the scientific community: Exploring the role of retractions in biomedicine. *Res. Policy*, 41(2):276–290, 2012. ISSN 00487333. doi: 10.1016/j.respol.2011.11.001.

D. Goncalves, J. Libgober, and J. Willis. Learning versus Unlearning: An Experiment on Retractions. 2021.

A. C. Gourieroux, A. Monfort, and A. Trognon. Pseudo Maximum Likelihood Methods: Theory. *Econometrica*, 52(3):681–700, 1984.

F. Hesselmann, V. Graf, M. Schmidt, and M. Reinhart. The visibility of scientific misconduct: A review of the literature on retracted journal articles. *Curr. Sociol. Rev.*, 65(6):814–845, 2017. ISSN 14617064. doi: 10.1177/0011392116663807.

A. Ivanova, M. S. Schäfer, I. Schlichting, and A. Schmidt. Is There a Medialization of Climate Science? Results From a Survey of German Climate Scientists. *Sci. Commun.*, 35(5):626–653, 2013. ISSN 10755470. doi: 10.1177/1075547012475226.

G. Z. Jin, B. Jones, S. F. Lu, and B. Uzzi. The reverse matthew effect: Consequences of retraction in scientific teams. *Rev. Econ. Stat.*, 101(3):492–506, 2019. ISSN 15309142. doi: 10.1162/rest_a_00780.

D. M. J. Lazer, M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, M. Schudson, S. A. Sloman, C. R. Sunstein, E. A.

Thorson, D. J. Watts, and J. L. Zittrain. The science of fake news. *Science (80-.).*, 359(6380): 1094–1096, mar 2018. ISSN 0036-8075. doi: 10.1126/science.aao2998. URL `https://www.science.org/doi/10.1126/science.aao2998`.

S. Lewandowsky, U. K. Ecker, C. M. Seifert, N. Schwarz, and J. Cook. Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychol. Sci. Public Interes. Suppl.*, 13(3):106–131, 2012. ISSN 15291006. doi: 10.1177/1529100612451018.

S. F. Lu, G. Z. Jin, B. Uzzi, and B. Jones. The Retraction Penalty: Evidence from the Web of Science. *Sci. Rep.*, 3:1–5, 2013. ISSN 20452322. doi: 10.1038/srep03146.

W. Matsuyama, H. Mitsuyama, M. Watanabe, K.-i. Oonakahara, I. Higashimoto, M. Osame, and K. Arimura. RETRACTED: Effects of Omega-3 Polyunsaturated Fatty Acids on Inflammatory Markers in COPD. *Chest*, 128(6):3817–3827, dec 2005. ISSN 00123692. doi: 10.1378/chest.128.6.3817. URL `https://linkinghub.elsevier.com/retrieve/pii/S0012369215496230`.

M. R. Mehra, S. S. Desai, S. Kuy, T. D. Henry, and A. N. Patel. RETRACTED: Cardiovascular Disease, Drug Therapy, and Mortality in Covid-19. *N. Engl. J. Med.*, 382(25):e102, 2020a. ISSN 0028-4793. doi: 10.1056/nejmoa2007621.

M. R. Mehra, S. S. Desai, F. Ruschitzka, and A. N. Patel. RETRACTED: Hydroxychloroquine or chloroquine with or without a macrolide for treatment of COVID-19: a multinational registry analysis. *Lancet*, 2020b. ISSN 1474547X. doi: 10.1016/S0140-6736(20)31180-6. URL `http://dx.doi.org/10.1016/S0140-6736(20)31180-6`.

P. Mongeon and V. Larivière. Costly collaborations: The impact of scientific fraud on co-authors' careers. *J. Assoc. Inf. Sci. Technol.*, 67(3):535–542, 2016. ISSN 23301643. doi: 10.1002/asi.23421.

D. P. Phillips, E. J. Kanter, B. Bednarczyk, and P. L. Tastad. Importance of the Lay Press in the Transmission of Medical Knowledge to the Scientific Community. *N. Engl. J. Med.*, 325 (16):1180–1183, oct 1991. ISSN 0028-4793. doi: 10.1056/NEJM199110173251620. URL `http://www.nejm.org/doi/abs/10.1056/NEJM199110173251620`.

C. Piller. Disgraced COVID-19 studies are still routinely cited. *Science (80-.).*, 371(6527):331–332, jan 2021. ISSN 0036-8075. doi: 10.1126/science.371.6527.331. URL `https://www.sciencemag.org/lookup/doi/10.1126/science.371.6527.331`.

M. F. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980. ISSN 00330337. doi: 10.1108/eb046814.

A. Prat and D. Strömberg. The Political Economy of Mass Media. In D. Acemoglu, M. Arellano, and E. Dekel, editors, *Adv. Econ. Econom. Tenth World Congr.*, pages 135–187. Cambridge University Press, Cambridge, 2013. ISBN 9781139060028. doi: 10.1017/CBO9781139060028.004. URL `https://www.cambridge.org/core/product/identifier/CBO9781139060028A013/type/book{_}part`.

D. Rosati. Citations are not opinions: a corpus linguistics approach to understanding how citations are made. apr 2021. URL `http://arxiv.org/abs/2104.08087`.

D. Sarathchandra and A. M. McCright. The Effects of Media Coverage of Scientific Retractions on Risk Perceptions. *SAGE Open*, 7(2), 2017. ISSN 21582440. doi: 10.1177/2158244017709324.

J. Schneider, D. Ye, A. M. Hill, and A. S. Whitehorn. *Continued post-retraction citation of a fraudu-*

lent clinical trial report, 11Â years after it was retracted for falsifying data, volume 125. 2020. ISBN 0123456789. doi: 10.1007/s11192-020-03631-1.

S. Serghiou, R. M. Marton, and J. P. Ioannidis. Media and social media attention to retracted articles according to Altmetric. *PLoS One*, 16(5 May 2021):1–19, 2021. ISSN 19326203. doi: 10.1371/journal.pone.0248625. URL `http://dx.doi.org/10.1371/journal.pone.0248625`.

M. Serra-Garcia and U. Gneezy. Nonreplicable publications are cited more than replicable ones. *Sci. Adv.*, 7(21):1–23, may 2021. ISSN 2375-2548. doi: 10.1126/sciadv.abd1705. URL `https://www.science.org/doi/10.1126/sciadv.abd1705`.

Y. Sugawara, T. Tanimoto, S. Miyagawa, M. Murakami, A. Tsuya, A. Tanaka, M. Kami, and H. Narimatsu. Scientific misconduct and social media: Role of twitter in the stimulus triggered acquisition of pluripotency cells scandal. *J. Med. Internet Res.*, 19(2):1–10, 2017. ISSN 14388871. doi: 10.2196/jmir.6706.

P. Sumner, S. Vivian-Griffiths, J. Boivin, A. Williams, C. A. Venetis, A. Davies, J. Ogden, L. Whelan, B. Hughes, B. Dalton, F. Boy, and C. D. Chambers. The association between exaggeration in health related science news and academic press releases: Retrospective observational study. *BMJ*, 349(December):1–8, 2014. ISSN 17561833. doi: 10.1136/bmj.g7015. URL `http://dx.doi.org/doi:10.1136/bmj.g7015`.

P. Sumner, S. Vivian-Griffiths, J. Boivin, A. Williams, L. Bott, R. Adams, C. A. Venetis, L. Whelan, B. Hughes, and C. D. Chambers. Exaggerations and caveats in press releases and health-related science news. *PLoS One*, 11(12):1–15, 2016. ISSN 19326203. doi: 10.1371/journal.pone.0168217.

I. Tahamtan, A. Safipour Afshar, and K. Ahamdzadeh. Factors affecting number of citations: a comprehensive review of the literature. *Scientometrics*, 107(3):1195–1225, 2016. ISSN 15882861. doi: 10.1007/s11192-016-1889-2.

The Economist. Zombie research haunts academic literature long after its supposed demise, 2021. URL `https://www.economist.com/graphic-detail/2021/06/26/zombie-research-haunts-academic-literature-long-after-its-supposed-demise`.

P. Weingart. Science and the media. *Res. Policy*, 27(8):869–879, dec 1998. ISSN 00487333. doi: 10.1016/S0048-7333(98)00096-1. URL `https://doi.org/10.1016/S0048-7333(98)00096-1https://linkinghub.elsevier.com/retrieve/pii/S0048733398000961`.

W. P. Whitely. The scientific community's response to evidence of fraudulent publication. The Robert Slutsky case. *JAMA J. Am. Med. Assoc.*, 272(2):170–173, 1994. ISSN 00987484. doi: 10.1001/jama.272.2.170.

K. Wohlrabe and C. Bürgi. Do working papers increase journal citations? Evidence from the top 5 journals in economics. *Appl. Econ. Lett.*, 28(17):1531–1535, 2021. ISSN 14664291. doi: 10.1080/13504851.2020.1855303. URL `https://doi.org/10.1080/13504851.2020.1855303`.

L. Ziegler. What is the Media Impact of Research in Economics ? 2021.

Appendix

Figures

Figure A.1: Control quality: citations.

Panel A: AD Panel B: ED Panel C: Cum. citations

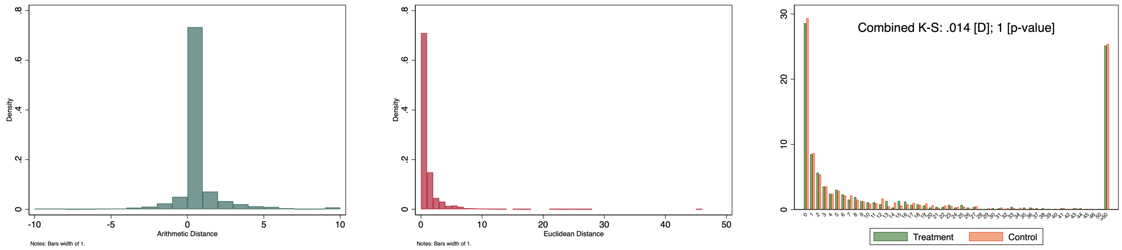

Note: All panels refer to pre-retraction measures. The year of retraction in that of the corresponding treated paper. Panel A shows the distribution of arithmetic distance (AD), panel B shows the distribution of Euclidean distance (ED), and panel C shows the distribution of cumulative citations from publication to the year before retraction and display the result of the Kolmogorov-Smirnov test of the equality of distributions between treatment and control group.

Figure A.2: Control quality: early mentions.

Panel A: Newspaper early mentions Panel B: Blog early mentions

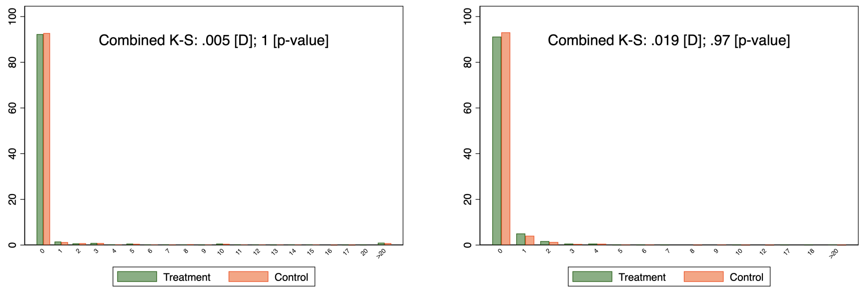

Note: Panels display the distribution of online mentions within two weeks from publication in newspapers (panel A) and blogs (panel B) across treated (green) and control (orange) papers. Both graph report the result of the Kolmogorov-Smirnov test of the equality of distributions across groups.

Figure A.3: Newspaper and blog mentions of selected control articles.

Panel A: News mentions (N=235 controls) Panel B: Blog mentions (N=216 controls)

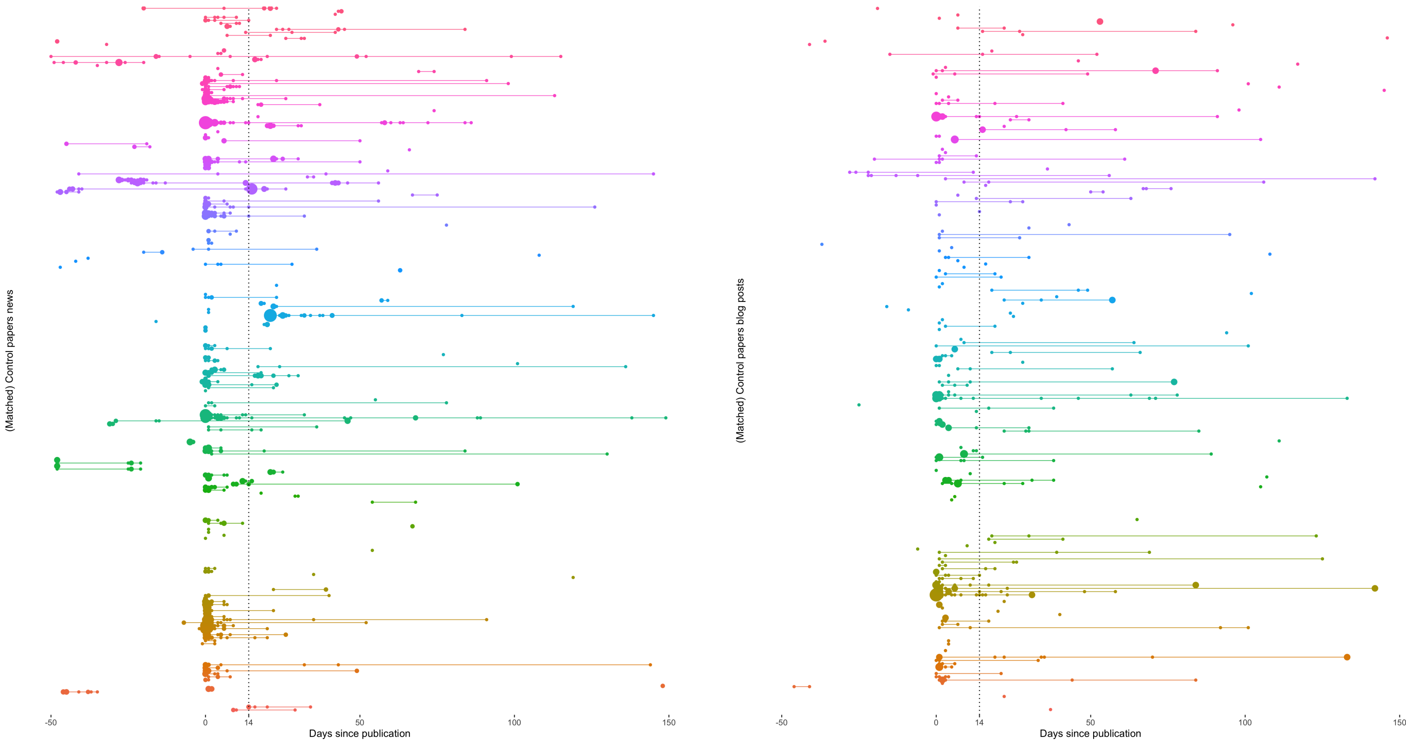

Note: Each line connects the first to the last mention of a single research article on either newspapers (Panel A) or blogs (Panel B) within the considered time window. Dots represent the number of mentions at a certain point in time. The source of publication date is *Altmetric*.

Figure A.4: Dynamics of retracted papers penalty

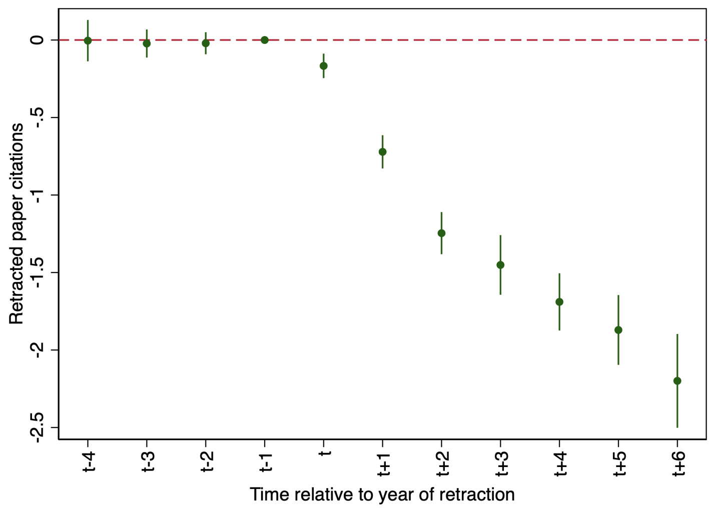

Note: Estimates replicate the model in Table A.2 column (2) but replacing the *Post* indicator with a full set of dummies for each year relative to the retraction ($t-1$ excluded). The coefficient displayed are that of the interaction between time dummies and a treatment indicator while the vertical lines represent 95% CI.

Tables

Table A.1: Altscore weights

News	8
Blog	5
Policy document (per source)	3
Patent	3
Wikipedia	3
Twitter (tweets and retweets)	1
Peer review (Publons, Pubpeer)	1
Weibo (not trackable since 2015, but historical data kept)	1
Google+ (not trackable since 2019, but historical data kept)	1
F1000	1
Syllabi (Open Syllabus)	1
LinkedIn (not trackable since 2014, but historical data kept)	0.5
Facebook (only a curated list of public Pages)	0.25
Reddit	0.25
Pinterest (not trackable since 2013, but historical data kept)	0.25
Q&A (Stack Overflow)	0.25
Youtube	0.25
Number of Mendeley readers	0
Number of Dimensions and Web of Science citations	0

Table A.2: Retracted papers penalty

	Poisson	Poisson	OLS
	Citations	Citations	IHS(Citations)
Post	0.122***	0.115***	0.150***
	(0.024)	(0.025)	(0.025)
Post * Treatment	-1.067***	-1.064***	-0.830***
	(0.062)	(0.060)	(0.030)
Article FE	Y	Y	Y
Age FE	Y	Y	Y
Year FE	N	Y	Y
Pseudo R2 / R2	0.708	0.708	0.772
N	15438	15438	16679
N clusters	966	966	979
N full	16711	16711	16711

Note: First two columns show estimates of pseudo Poisson specifications while third column shows OLS estimation with IHS transformed dependent variable. The dependent variable is the total number of citations (exclusive of self-citations) received by each article in a particular year. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.3: Retracted papers penalty with early visibility

	(1) Citations	(2) Citations	(3) Citations
Post * Treatment	-0.959***	-0.983***	-0.977***
	(0.059)	(0.058)	(0.058)
Post * Treatment * Early visibility	-0.449***		
	(0.158)		
Post * Treatment * Early blog visibility		-0.396**	
		(0.185)	
Post * Treatment * Early news visibility			-0.418***
			(0.149)
Article FE	Y	Y	Y
Age FE	Y	Y	Y
Year FE	Y	Y	Y
Pseudo R2	0.709	0.709	0.709
N	15438	15438	15438
N clusters	966	966	966
N full	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Early visibility is an indicator for papers with at least one mention (in newspapers and/or blogs) within two weeks from publication. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.4: Retracted papers penalty with any media coverage

	(1) Citations	(2) Citations	(3) Citations	(4) Citations
Post * Treatment	-0.831***	-0.798***	-0.944***	-0.840***
	(0.083)	(0.072)	(0.059)	(0.077)
Post * Treatment * Any social media	-0.325***			
	(0.119)			
Post * Treatment * Any news-blog		-0.434***		
		(0.119)		
Post * Treatment * Any news			-0.377***	
			(0.127)	
Post * Treatment * Any blog				-0.392***
				(0.129)
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.709	0.709	0.709	0.709
N	15438	15438	15438	15438
N clusters	966	966	966	966
N full	16711	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Any socialmedia/news/blog is an indicator for papers with at least one overall mention in any of the indicated outlets. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.5: Retracted papers penalty and attention score

	(1) Citations	(2) Citations	(3) Citations	(4) Citations
Post * Treatment	-0.841***	-0.815***	-0.921***	
	(0.090)	(0.068)	(0.051)	
Post * Treatment * Altscore >p50	-0.283**			
	(0.117)			
Post * Treatment * Altscore >p75		-0.433***		
		(0.118)		
Post * Treatment * Altscore >p90			-0.488***	
			(0.150)	
Post * Treatment * Altscore 3rd quintile				-0.732***
				(0.091)
Post * Treatment * Altscore 4th quintile				-0.918***
				(0.083)
Post * Treatment * Altscore 5th quintile				-1.267***
				(0.088)
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.709	0.709	0.709	0.709
N	15438	15438	15438	15438
N clusters	966	966	966	966
N full	16711	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Altscore is an aggregate measure of weighted online mentions. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.6: Retracted papers penalty with early visibility ($Post\ t \geq 0$)

VARIABLES	(1) Citations	(2) Citations	(3) Citations
Post * Treatment	-0.653***	-0.668***	-0.657***
	(0.061)	(0.061)	(0.060)
Post * Treatment * Early visibility	-0.327**		
	(0.138)		
Post* Treatment * Early blog visibility		-0.268	
		(0.170)	
Post * Treatment * Early news visibility			-0.347**
			(0.139)
Article FE	Y	Y	Y
Age FE	Y	Y	Y
Year FE	Y	Y	Y
Pseudo R2	0.701	0.701	0.701
N	15438	15438	15438
N clusters	966	966	966
N full	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for the year of retraction and all subsequent years. Early visibility is an indicator for papers with at least one mention (in newspapers and/or blogs) within two weeks from publication. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.7: Retracted papers penalty with any media coverage (*Post* $t \geq 0$)

VARIABLES	(1) Citations	(2) Citations	(3) Citations	(4) Citations
Post * Treatment	-0.465***	-0.465***	-0.608***	-0.527***
	(0.105)	(0.092)	(0.065)	(0.088)
Post * Treatment * Any social media	-0.346***			
	(0.128)			
Post * Treatment * Any news-blog		-0.414***		
		(0.126)		
Post * Treatment * Any news			-0.355***	
			(0.115)	
Post * Treatment * Any blog				-0.339***
				(0.129)
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.701	0.702	0.701	0.701
N	15438	15438	15438	15438
N clusters	966	966	966	966
N full	16711	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for the year of retraction and all subsequent years. Any socialmedia/news/blog is an indicator for papers with at least one overall mention in any of the indicated outlets. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.8: Retracted papers penalty and attention score ($Post\ t \geq 0$)

VARIABLES	(1) Citations	(2) Citations	(3) Citations	(4) Citations
Post * Treatment	-0.452***	-0.467***	-0.572***	
	(0.118)	(0.083)	(0.058)	
Post * Treatment * Altscore >p50	-0.339**			
	(0.133)			
Post * Treatment * Altscore >p75		-0.439***		
		(0.118)		
Post * Treatment * Altscore >p90			-0.554***	
			(0.127)	
Post * Treatment * Altscore 3rd quintile				-0.521***
				(0.089)
Post * Treatment * Altscore 4th quintile				-0.607***
				(0.082)
Post * Treatment * Altscore 5th quintile				-0.919***
				(0.072)
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.701	0.702	0.702	0.702
N	15438	15438	15438	15438
N clusters	966	966	966	966
N full	16711	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for the year of retraction and all subsequent years. Altscore is an aggregate measure of weighted online mentions. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.9: Retracted papers penalty with early visibility (actively cited papers)

	(1) Citations	(2) Citations	(3) Citations
Post * Treatment	-1.036***	-1.078***	-1.064***
	(0.089)	(0.086)	(0.090)
Post * Treatment * Early visibility	-0.428**		
	(0.198)		
Post * Treatment * Early blog visibility		-0.362	
		(0.238)	
Post * Treatment * Early news visibility			-0.399**
			(0.165)
Article FE	Y	Y	Y
Age FE	Y	Y	Y
Year FE	Y	Y	Y
Pseudo R2	0.733	0.733	0.733
N	7308	7308	7308
N clusters	466	466	466
N full	7662	7662	7662

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Early visibility is an indicator for papers with at least one mention (in newspapers and/or blogs) within two weeks from publication. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.10: Retracted papers penalty with any media coverage (actively cited papers)

	(1) Citations	(2) Citations	(3) Citations	(4) Citations
Post * Treatment	-0.739***	-0.785***	-1.010***	-0.892***
	(0.124)	(0.113)	(0.093)	(0.113)
Post * Treatment * Any social media	-0.577***			
	(0.162)			
Post * Treatment * Any news-blog		-0.618***		
		(0.162)		
Post * Treatment * Any news			-0.387***	
			(0.139)	
Post * Treatment * Any blog				-0.494***
				(0.171)
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.733	0.734	0.733	0.733
N	7308	7308	7308	7308
N clusters	466	466	466	466
N full	7662	7662	7662	7662

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Any socialmedia/news/blog is an indicator for papers with at least one overall mention in any of the indicated outlets. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.11: Retracted papers penalty and attention score (actively cited papers)

	(1) Citations	(2) Citations	(3) Citations	(4) Citations
Post * Treatment	-0.704***	-0.805***	-0.967***	
	(0.132)	(0.106)	(0.080)	
Post * Treatment * Altscore >p50	-0.583***			
	(0.163)			
Post * Treatment * Altscore >p75		-0.590***		
		(0.157)		
Post * Treatment * Altscore >p90			-0.544***	
			(0.175)	
Post * Treatment * Altscore 3rd quintile				-0.769***
				(0.149)
Post * Treatment * Altscore 4th quintile				-1.123***
				(0.119)
Post * Treatment * Altscore 5th quintile				-1.317***
				(0.102)
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.733	0.733	0.734	0.733
N	7308	7308	7308	7308
N clusters	466	466	466	466
N full	7662	7662	7662	7662

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Altscore is an aggregate measure of weighted online mentions. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.12: Selected summary statistics: title ngrams

Most frequent (selected) ngrams					Relevant selected ngrams				
	Mean	Sd	Min	Max		Mean	Sd	Min	Max
Treatment	0.0477	0.213	0	1	# of adult	0.0070	0.084	0	2
Media	0.0945	0.293	0	1	# of algorithm	0.0113	0.107	0	2
Citations $year_p$	0.901	3.217	0	249	# of brain	0.009	0.097	0	1
Total number of words	14.28	5.047	1	54	# of climat	0.005	0.071	0	2
# of base	0.0773	0.272	0	2	# of commun	0.0080	0.091	0	2
# of effect	0.0666	0.252	0	2	# of composit	0.0201	0.142	0	2
# of studi	0.0543	0.228	0	2	# of disord	0.0064	0.083	0	2
# of model	0.0517	0.225	0	2	# of earli	0.0075	0.086	0	1
# of analysi	0.0453	0.209	0	2	# of genom	0.0100	0.101	0	2
# of system	0.0380	0.195	0	2	# of global	0.0065	0.082	0	2
# of induc	0.0314	0.176	0	2	# of graphen	0.0093	0.101	0	3
# of imag	0.0291	0.172	0	3	# of meta_analysi	0.0050	0.070	0	1
# of human	0.0270	0.165	0	2	# of model	0.0517	0.225	0	2
# of perform	0.0256	0.159	0	2	# of network_ETX	0.0082	0.090	0	1
# of mechan	0.0252	0.158	0	2	# of neuron	0.0052	0.075	0	2
# of properti	0.0244	0.155	0	2	# of reveal	0.0092	0.096	0	1
# of oxid	0.0242	0.163	0	3	# of risk	0.0154	0.127	0	3
# of enhanc	0.0238	0.153	0	2	# of stem	0.0095	0.100	0	2
# of regul	0.0238	0.154	0	2	# of STX_structur	0.0057	0.076	0	1
# of associ	0.0236	0.155	0	2	# of trial	0.0108	0.104	0	2
# of respons	0.0233	0.153	0	2	# of vitro	0.0064	0.080	0	1

Note: N-grams represent the number of times the selected espression appears in the title of a research article. All n-grams in the table were selected by one of the lasso procedures. N=20755.

Table A.13: Likelihood of retraction and media coverage (Logit)

	Retraction									
	Logit		EBIC		AICC		CV		Rigourous	
Media coverage	0.009**	0.018***	0.014***	0.020***	0.015***	0.021***	0.015***	0.021***	0.014***	0.020***
	(0.004)	(0.006)	(0.004)	(0.006)	(0.004)	(0.006)	(0.004)	(0.006)	(0.004)	(0.006)
Predicted media			-0.089***	-0.085**	-0.054***	-0.054**	-0.058***	-0.058***	-0.095**	-0.091**
			(0.033)	(0.038)	(0.016)	(0.021)	(0.017)	(0.022)	(0.037)	(0.043)
Publication year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Journal FE	N	Y	N	Y	N	Y	N	Y	N	Y
N	20755	20650	20755	20650	20755	20650	20755	20650	20755	20650

Note: Estimates from Logit equation. The dependent variable *Retraction* is an indicator for whether a paper was retracted. *Media coverage* is an indicator for whether a paper attracted online coverage at publication. *Predicted media* is media coverage as predicted from the respective *lassologit* procedures. Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.14: Likelihood of retraction and media coverage (selection within subjects, publication years and excluding retractions)

	Retraction							
	EBIC		AICC		CV		Rigourous	
Media coverage	0.016***	0.021***	0.018***	0.022***	0.018***	0.022***	0.014***	0.020***
	(0.005)	(0.007)	(0.005)	(0.007)	(0.005)	(0.007)	(0.005)	(0.007)
Predicted media	-0.056***	-0.071**	-0.061***	-0.074***	-0.062***	-0.074***	-0.047***	-0.084**
	(0.017)	(0.028)	(0.015)	(0.023)	(0.015)	(0.023)	(0.014)	(0.035)
Publication year FE	Y	Y	Y	Y	Y	Y	Y	Y
Journal FE	N	Y	N	Y	N	Y	N	Y
N	20393	20393	20393	20393	20393	20393	20393	20393
N clusters	988	988	988	988	988	988	988	988

Note: Estimates from OLS regression. The dependent variable *Retraction* is an indicator for whether a paper was retracted. *Media coverage* is an indicator for whether a paper attracted online coverage at publication. *Predicted media* is media coverage as predicted from the respective *lasso* procedures. Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.15: Months to retraction and Journal-year average visibility

	(1) Time to Retract	(2) Time to Retract	(3) Time to Retract	(4) Time to Retract	(5) Time to Retract	(6) Time to Retract
Altscore	-3.579***					
	(0.565)					
Sh. Blog		-3.376				
		(2.413)				
Blog count		-1.662				
		(1.850)				
Sh. news			-6.287***			
			(1.937)			
News count			2.105			
			(1.384)			
Sh. Tweets				-3.127***		
				(1.117)		
Tweets count				-1.688***		
				(0.645)		
Sh. early blog					-3.331*	
					(1.995)	
Early blog count					-0.423	
					(1.675)	
Sh. early news						-5.453***
						(1.666)
Early news count						1.931
						(1.262)
Observations	967	961	962	968	962	967
R-squared	0.455	0.468	0.468	0.460	0.459	0.455
Publication year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Mean	24.21	24.21	24.21	24.21	24.21	24.21

Note: The dependent variable is the time intercurring between an article publication and its retraction, expressed in months. Covariates represents different measures of journal visibility, measured as the average of non-retracted papers that appear in the same yournal and year of the retracted one. All covariates are standardized and outliers trimmed. Controls include the number of non-retracted articles within same journal and year of the treated, the average Euclidean distance of those from the retracted paper, and the level of (non-self) cumulative citations of the retracted paper before retraction. Publication year fixed effects are included. Journal clustered standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.16: Loss in citation and Journal-year average visibility

	(1) DID (-1,1)	(2) DID (-1,1)	(3) DID (-2,2)	(4) DID (-2,2)	(5) DID (-4,4)	(6) DID (-4,4)	(7) DID (-4,6)	(8) DID (-4,6)
Altscore	-1.362***	-1.111***	-1.699***	-1.327***	-1.830***	-1.399***	-1.877***	-1.439***
	(0.187)	(0.221)	(0.201)	(0.202)	(0.224)	(0.221)	(0.236)	(0.232)
Observations	840	840	840	840	840	840	840	840
R-squared	0.120	0.171	0.165	0.230	0.158	0.236	0.152	0.231
Sh. Blog	-1.698**	-1.643**	-1.598**	-1.440**	-1.746**	-1.548**	-1.658**	-1.449**
	(0.738)	(0.686)	(0.740)	(0.677)	(0.767)	(0.687)	(0.764)	(0.686)
Blog count	0.078	0.268	-0.327	-0.088	-0.334	-0.066	-0.467	-0.198
	(0.651)	(0.627)	(0.601)	(0.562)	(0.614)	(0.562)	(0.615)	(0.567)
Observations	831	831	831	831	831	831	831	831
R-squared	0.139	0.188	0.182	0.241	0.175	0.246	0.167	0.239
Sh. News	-1.611***	-1.365***	-1.795***	-1.325**	-1.952***	-1.384**	-1.995***	-1.406**
	(0.519)	(0.501)	(0.582)	(0.559)	(0.633)	(0.599)	(0.650)	(0.611)
News count	0.129	0.107	-0.018	-0.128	-0.051	-0.194	-0.059	-0.214
	(0.446)	(0.421)	(0.480)	(0.461)	(0.515)	(0.489)	(0.525)	(0.495)
Observations	835	835	835	835	835	835	835	835
R-squared	0.125	0.173	0.169	0.228	0.167	0.238	0.161	0.233
Sh. Tweets	-0.650**	-0.540**	-0.764***	-0.558**	-0.856***	-0.601**	-0.894***	-0.627**
	(0.268)	(0.262)	(0.275)	(0.257)	(0.295)	(0.270)	(0.298)	(0.273)
Tweets count	-0.965***	-0.788***	-1.185***	-0.944***	-1.233***	-0.966***	-1.254***	-0.986***
	(0.222)	(0.238)	(0.251)	(0.235)	(0.286)	(0.261)	(0.302)	(0.275)
Observations	842	842	842	842	842	842	842	842
R-squared	0.123	0.173	0.164	0.229	0.156	0.234	0.149	0.228
Publication year FE	Y	Y	Y	Y	Y	Y	Y	Y
Age at retraction FE	N	Y	N	Y	N	Y	N	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Mean	-3.149	-3.149	-3.707	-3.707	-4.274	-4.274	-4.158	-4.158

Note: The dependent variable is the individual loss in citations obtained comparing each retracted paper to its selected controls for different pre- and post- time windows. Covariates represents different measures of journal visibility, measured as the average of non-retracted papers that appear in the same yournal and year of the retracted one. All covariates are standardized and outliers trimmed. Controls include the number of non-retracted articles within same journal and year of the treated, the average Euclidean distance of those from the retracted paper, and the level of (non-self) cumulative citations of the retracted paper before retraction. Publication year fixed effects are included. Fixed effects for age of the article at retraction are added in even comuns. Journal clustered standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.17: Retracted papers penalty and early visibility by discipline

	(1)	(2)	(3) Hard sciences	(4) Social sciences
	Citations	Citations	Citations	Citations
Post * Treatment * Early visibility	0.592	0.364	-0.520***	0.366
	(0.396)	(0.262)	(0.163)	(0.264)
Post * Treatment * Early visibility * Business/Technology	-0.155			
	(0.522)			
Post * Treatment * Early visibility * Life sciences	-1.169**			
	(0.458)			
Post * Treatment * Early visibility * Environment	-0.917**			
	(0.457)			
Post * Treatment * Early visibility * Health	-1.195**			
	(0.529)			
Post * Treatment * Early visibility * Physics	-0.938**			
	(0.455)			
Post * Treatment * Early visibility * Hard sciences		-0.887***		
		(0.309)		
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.711	0.710	0.718	0.595
N	15399	15438	12980	2419
N clusters	964	966	798	166
N full	16672	16711	13837	2835

Note: Estimates derive from pseudo Poisson specifications. Hard sciences: life sciences, environment, health and physical sciences. Social sciences: business and technology, humanities, other social sciences. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Early visibility is an indicator for papers with at least one mention (in newspapers and/or blogs) within two weeks from publication. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.18: Retracted papers penalty and attention score by discipline

	(1)	(2)	(3) Hard sciences	(4) Social sciences
	Citations	Citations	Citations	Citations
Post * Treatment * Altscore >p50	1.333**	0.167	-0.321**	0.150
	(0.622)	(0.254)	(0.127)	(0.264)
Post * Treatment * Altscore >p50 * Business/Technology	-1.307*			
	(0.680)			
Post * Treatment * Altscore >p50 * Life sciences	-1.954***			
	(0.650)			
Post * Treatment * Altscore >p50 * Environment	-1.197*			
	(0.717)			
Post * Treatment * Altscore >p50 * Health	-0.948			
	(0.692)			
Post * Treatment * Altscore >p50 * Physics	-1.485**			
	(0.649)			
Post * Treatment * Altscore >p50 * Hard sciences		-0.494*		
		(0.283)		
Article FE	Y	Y	Y	Y
Age FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Pseudo R2	0.710	0.709	0.718	0.595
N	15399	15438	12980	2419
N clusters	964	966	798	166
N full	16672	16711	13837	2835

Note: Estimates derive from pseudo Poisson specifications. Hard sciences: life sciences, environment, health and physical sciences. Social sciences: business and technology, humanities, other social sciences. The dependent variable is the total number of citations (exclusive of self-citations) received by each paper in a particular year. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Altscore is an aggregate measure of weighted online mentions. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citations). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.19: Citation statements and visibility

	(1) Cit. statements	(2) Cit. statements	(3) Mentioning	(4) Mentioning	(5) Contrasting	(6) Contrasting	(7) Supporting	(8) Supporting
Post	0.033	0.038	0.034	0.040	-0.108	-0.132	0.035	0.031
	(0.038)	(0.040)	(0.038)	(0.040)	(0.236)	(0.263)	(0.100)	(0.109)
Post * Treatment	-1.215***	-1.165***	-1.231***	-1.178***	-1.147***	-1.184***	-1.055***	-1.061***
	(0.096)	(0.093)	(0.095)	(0.093)	(0.339)	(0.348)	(0.197)	(0.209)
Post * Early visibility	0.245***	0.171***	0.238***	0.167**	0.738**	0.745**	0.417***	0.306*
	(0.063)	(0.066)	(0.065)	(0.068)	(0.317)	(0.347)	(0.158)	(0.170)
Post * Treatment * Early visibility	-0.411**	-0.486***	-0.396**	-0.478***	-0.595	-0.416	-0.456	-0.455
	(0.164)	(0.170)	(0.167)	(0.173)	(0.684)	(0.730)	(0.414)	(0.440)
Self cit. excluded	N	Y	N	Y	N	Y	N	Y
Article FE	Y	Y	Y	Y	Y	Y	Y	Y
Age FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Pseudo R2	0.713	0.717	0.711	0.715	0.138	0.138	0.267	0.254
N	14594	14158	14536	14089	1701	1421	6370	5761
N clusters	957	948	956	946	201	170	567	521
N full	16711	16711	16711	16711	16711	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations statements received by each paper in a particular year, even columns exclude self citations. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Early visibility is an indicator for papers with at least one mention (in newspapers and/or blogs) within two weeks from publication. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citation statements). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.20: Citation statements and attention score

	(1) Cit. statements	(2) Cit. statements	(3) Mentioning	(4) Mentioning	(5) Contrasting	(6) Contrasting	(7) Supporting	(8) Supporting
Post	-0.023	0.024	-0.022	0.024	0.050	0.352	0.045	0.115
	(0.052)	(0.055)	(0.053)	(0.056)	(0.344)	(0.360)	(0.131)	(0.142)
Post * Treatment	-1.001***	-0.941***	-0.992***	-0.923***	-1.696***	-1.972***	-1.218***	-1.345***
	(0.121)	(0.119)	(0.121)	(0.119)	(0.634)	(0.681)	(0.351)	(0.373)
Post * Altscore >p50	0.173***	0.090	0.170***	0.091	0.048	-0.367	0.149	0.009
	(0.060)	(0.063)	(0.060)	(0.064)	(0.350)	(0.368)	(0.151)	(0.159)
Post * Treatment * Altscore >p50	-0.415***	-0.441***	-0.441***	-0.477***	0.480	0.808	0.030	0.190
	(0.155)	(0.153)	(0.155)	(0.153)	(0.704)	(0.756)	(0.403)	(0.428)
Self cit. excluded	N	Y	N	Y	N	Y	N	Y
Article FE	Y	Y	Y	Y	Y	Y	Y	Y
Age FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Pseudo R2	0.713	0.717	0.711	0.715	0.136	0.136	0.266	0.253
N	14594	14158	14536	14089	1701	1421	6370	5761
N clusters	957	948	956	946	201	170	567	521
N full	16711	16711	16711	16711	16711	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations statements received by each paper in a particular year, even columns exclude self citations. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Altscore is an aggregate measure of weighted online mentions. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citation statements). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.

Table A.21: Citation statements and attention score extremes

	(1) Cit. statements	(2) Cit. statements	(3) Mentioning	(4) Mentioning	(5) Contrasting	(6) Contrasting	(7) Supporting	(8) Supporting
Post	0.043	0.074	0.046	0.074	0.262	0.567	0.070	0.177
	(0.056)	(0.059)	(0.056)	(0.060)	(0.373)	(0.355)	(0.146)	(0.157)
Post * Treatment	-1.103***	-1.002***	-1.097***	-0.984***	-1.940**	-2.203***	-1.226***	-1.385***
	(0.156)	(0.153)	(0.155)	(0.150)	(0.769)	(0.820)	(0.467)	(0.484)
Post * Treatment * Altscore 3rd qntl	0.428**	0.287	0.438**	0.286	1.369	1.523	0.203	0.315
	(0.211)	(0.217)	(0.212)	(0.218)	(1.254)	(1.342)	(0.654)	(0.684)
Post * Treatment * Altscore 4th qntl	-0.047	-0.105	-0.087	-0.160	0.249	0.468	0.476	0.743
	(0.205)	(0.205)	(0.204)	(0.204)	(1.288)	(1.338)	(0.546)	(0.568)
Post * Treatment * Altscore 5th qntl	-0.456**	-0.510***	-0.472**	-0.538***	0.513	0.899	-0.224	-0.024
	(0.197)	(0.193)	(0.196)	(0.191)	(0.849)	(0.907)	(0.534)	(0.554)
Self cit. excluded	N	Y	N	Y	N	Y	N	Y
Age FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Pseudo R2	0.714	0.717	0.712	0.716	0.142	0.140	0.268	0.255
N	14594	14158	14536	14089	1701	1421	6370	5761
N clusters	957	948	956	946	201	170	567	521
N full	16711	16711	16711	16711	16711	16711	16711	16711

Note: Estimates derive from pseudo Poisson specifications. The dependent variable is the total number of citations statements received by each paper in a particular year, even columns exclude self citations. Treatment is an indicator for retracted papers. Post is an indicator for all years strictly after the year of retraction. Altscore is an aggregate measure of weighted online mentions. All models incorporate article fixed effects, a full suite of calendar-year effects and article age indicator variables. Using the following transformation $(1 - exp[\beta]) * 100 = x$ coefficients can be interpreted as elasticities (i.e. x% loss in yearly citation statements). Standard errors in parentheses, clustered around retraction cases. *p < 0.10, **p < 0.05, ***p < 0.01.