

NBER WORKING PAPER SERIES

THE REVERSE MATTHEW EFFECT:
CATASTROPHE AND CONSEQUENCE IN SCIENTIFIC TEAMS

Ginger Zhe Jin
Benjamin Jones
Susan Feng Lu
Brian Uzzi

Working Paper 19489
<http://www.nber.org/papers/w19489>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2013

We thank Alex Entz, Yiyang Liu, Huan Meng, and Ari Bellin for excellent research assistance and thank the University of Maryland, the Kellogg School of Management, the Northwestern University Institute on Complex Systems (NICO); the Army Research Laboratory under cooperative Agreement Number W911NF-09-2-0053 and Defense Advanced Research Projects Agency grant BAA-11-64, Social Media in Strategic Communication for financial support. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, the U.S. government, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Ginger Zhe Jin, Benjamin Jones, Susan Feng Lu, and Brian Uzzi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Reverse Matthew Effect: Catastrophe and Consequence in Scientific Teams
Ginger Zhe Jin, Benjamin Jones, Susan Feng Lu, and Brian Uzzi
NBER Working Paper No. 19489
October 2013, Revised May 2014
JEL No. J24,L15,L23,O3

ABSTRACT

What are the individual rewards to working in teams? This question extends across many production settings but is of long-standing interest in science and innovation, where the “Matthew Effect” suggests that eminent team members garner credit for great works at the expense of less eminent team members. In this paper, we study this question in reverse, examining highly negative events – article retractions. Using the Web of Science, we investigate how retractions affect citations to the authors’ prior publications. We find that the Matthew Effect works in reverse – namely, scientific misconduct imposes little citation penalty on eminent coauthors. By contrast, less eminent coauthors face substantial citation declines to their prior work, and especially when they are teamed with an eminent author. A simple Bayesian model is used to interpret the results. These findings suggest that a good reputation can have protective properties, but at the expense of those with less established reputations.

Ginger Zhe Jin
University of Maryland
Department of Economics
3115F Tydings Hall
College Park, MD 20742-7211
and NBER
jin@econ.umd.edu

Benjamin Jones
Northwestern University
Kellogg School of Management
Department of Management and Strategy
2001 Sheridan Road
Evanston, IL 60208
and NBER
bjones@kellogg.northwestern.edu

Susan Feng Lu
C3-219 Carol Simon Hall
Simon Graduate School of Business
University of Rochester
Rochester, NY, 14620
susan.lu@simon.rochester.edu

Brian Uzzi
Kellogg School of Management
2001 Sheridan Road
Evanston, IL 60208
uzzi@kellogg.northwestern.edu

1. Introduction

Team production is pervasive in modern economies, often related to the division of labor and benefits therein.¹ Yet team production raises challenges, including free riding during production and credit sharing concerns ex-post. In situations where the output of the individual is not directly observed, reputation may become a cornerstone not only in providing effort incentives but also in shaping how the community assigns credit across a team.

In a classic study, Robert K. Merton suggested the “Matthew Effect” as a fundamental issue in an important team production context, science (Merton 1968). Merton argued that more eminent coauthors tend to receive disproportionate credit for team-authored work (Merton 1968).² In Merton’s analysis, teamwork leads to a “rich get richer” phenomenon, where, faced with a great paper, the scientific community assumes that the more eminent coauthor was the key producer while less well-known coauthor(s) were subordinate contributors who deserve less credit. Arguably, such a credit assignment mechanism, if it operates, could have large effects on reputations, on the dynamics of individual careers, on incentives to work in teams, and on efficient matching of team members.

This paper considers a natural experiment to assess the individual consequences of working in teams. Our question, however, concerns not the rewards of “good” events, but rather consequences of catastrophes. Namely, we look at the effect of article retractions in team production settings and examine whether eminent coauthors attract or repel blame compared to less eminent coauthors. On the one hand, one might imagine that eminent authors receive disproportionate credit for the output, whether good or bad, as the presumed leader of the research enterprise. On the other hand, one

¹ See, e.g., classic observations in Bacon (1620) and Smith (1776) or modern analyses such as Becker and Murphy (1992), Hamilton et al. (2003), Jones (2010), and Mas and Moretti (2011).

² Merton coined the Matthew Effect after the biblical passage “For unto every one that hath shall be given, and he shall have abundance: but from him that hath not shall be taken even that which he hath” (Matthew 25: 29, King James Version).

may imagine that eminent authors have such established reputations that they escape blame for bad events, leaving any blame to accrue to junior coauthors. Thus we may imagine a “Reverse Matthew Effect”, which might be phrased as “To whom much has been given, much less will be taken away.”

In our empirical analysis, we collect retracted articles in the Web of Science where the retracted paper was authored in a team and where the authors have a single retraction event (that is, we do not look at extreme cases where an author is revealed to be a systematic fraud). We then investigate citation behavior to the prior publications of each author involved in the retracted work. To examine the effect of retraction, we match each of these prior publications (the treated papers) with a set of other publications (the control papers) that were published in the same field-year and received similar citations every year before the retraction event. This approach allows us to identify the effect of retraction via differences-in-differences estimation. This identification strategy builds from the observation that the content of prior work is unchanged, so that changes in citations to this work, compared to counterfactual control papers, reveal the effect of the retraction shock.³

Using standard measures of eminence from the science literature, we find three central results following retraction events. First, less established coauthors experience substantial citation declines to their prior work. Second, by contrast, eminent coauthors experience little or no citation consequences for their prior work. Third, less established authors are especially negatively affected in the presence of an eminent coauthor. This interaction effect suggests that eminence may act not only to protect oneself, but also to hurt others on one’s production team. These results persist across a variety of robustness checks.

Given these findings, and building from reasoning in Merton’s original Matthew Effect paper (Merton 1968), we further present a simple Bayesian model as a candidate explanation for the empirical results. In the model, the community attempts to infer

³ Using citations to prior scientific work to assess the effects of information shocks was pioneered as an identification strategy in Furman and Stern (2011).

each author's tendency to produce false science given different priors about each author and the possibility that anyone might make a mistake. Eminence is defined as a prior reputational state featuring precise beliefs that an author is a high quality type. In the presence of a retraction, the model shows that (1) being eminent helps you; (2) being eminent hurts your coauthors; and (3) eminence hurts a coauthor more the less established the coauthor is. The empirical results thus appear broadly consistent with a Bayesian inference problem, where the community assigns blame given priors over the individuals involved and their interactions.

The paper proceeds as follows. In Section 2 we review relevant literature and consider a range of qualitative theories that may bear on the response to negative events like a retraction. Section 3 presents the data and empirical strategy. Section 4 presents primary results. Section 5 develops a simple Bayesian model to provide a candidate explanation for the results and further discusses additional interpretations. Section 6 concludes.

2. Literature, Context, and Hypotheses

Team production is a ubiquitous feature of modern economies, where collaborative work is seen from restaurant kitchens to film production to satellite manufacturing. Teams have long been theorized to tap gains from specialization and the substantial productivity advantages therein (Smith 1776, Becker and Murphy 1992, Jones 2009). In practice, the U.S. Census currently indexes over 31,000 different occupational codes, and productivity gains from teamwork have been shown in settings from garment manufacturing (Hamilton, Nickerson and Owan 2003) to supermarket cashier services (Mas and Moretti 2011) to broad classes of scientific and inventive processes (Wuchty et al 2007; Jones et al 2008; Uzzi and Spiro 2005) where teams aggregate specialized knowledge (Jones 2009; Uzzi et al. 2013, Freeman et al. 2013).

Yet teamwork also raises agency problems. Indeed, the complementarities across individuals that can give teams their strength may also undermine their potential. For example, when individual contributions are not easily observed, it can be difficult for

outsiders to discern the effort or actions of individual team members. Team production can then be associated with free-riding problems and credit-sharing problems amidst other transaction costs associated with finding appropriate partners and ensuring efficient operation (e.g. Holmstrom 1982, Merton 1968, Hamilton, Nickerson, and Owan 2003, Cooper and Kagel, 2005). Thus, understanding team function in light of such challenges, especially given the ubiquity of teamwork and the productivity gains it can promise, is arguably a first-order question of broad application in modern economies.

Information challenges may be overcome through reputation and learning in many contexts, as suggested by large theoretical and empirical literatures. Reputation can be beneficial in establishing product quality, which may be difficult to accurately ascertain otherwise (Klein and Leffler 1981, Shapiro 1983). Generally, one can write a mapping

$$y = f(q, R) \tag{1}$$

where y is the realized demand for the output, which is increasing in both q , the quality of the output, and R , the reputation of the producer. If q is not fully observable to the buyer, then a good reputation may drive demand for the seller's products, as has been shown in settings from eBay transactions to medical services (Bajari and Hortacsu 2004, Pope 2009, Dranove, Ramanarayanan and Watanabe 2012). Sellers may then have natural incentives to obtain good reputations and avoid bad ones (Cabral and Hortacsu 2004, Jin and Leslie 2009, Johnson 2012).

Reputation, however, may have more complicated implications in settings of team production. Merton's "Matthew Effect" provides a canonical analysis (Merton 1968). Merton notes that the presence of a team member with a strongly positive reputation can enhance demand for the product (a research article in Merton's setting, where an eminent author attracts greater attention to the output) thus creating a positive spillover on other team members, especially junior researchers, by elevating attention to their work. This "communication" hypothesis is closely akin to the product market logic above, where a strong reputation, R , can enhance demand, y . On the other hand, and according to Merton's primary analysis, the presence of an eminent team

member may act to steal credit from the others, as the community infers that the eminent team member is responsible for the output. Thus, while partnering with a high-reputation teammate may enhance demand for the given output, it may also make it difficult for the less-established teammate to become established herself. In other words, Merton emphasizes a community inference problem, inverting the mapping (1), where an individual's reputation, R , becomes established through a series of outputs, y . Here an eminent team member may create a negative spillover on the other members, who may have contributed substantially to the production of y yet garner little credit or career advantage, as the community assumes the eminent team member was responsible for the success. This "credit" hypothesis may thus lead to a "rich get richer" phenomenon for which Merton coined the Matthew Effect. If this effect operates, it not only raises questions of fairness but may also create challenges in team production settings. For example, such a mechanism may slow career progress for young team members, perhaps dimming their interest in the career itself, as they struggle to establish independent reputations.⁴ More generally, ex-post credit considerations may disrupt efficient ex-ante formation of teams, as matches between individuals with appropriate complementary skills are now entangled with concerns over relative reputations.

Recent prior literature has examined Merton's communication hypothesis specifically in the setting of science and innovation. Simcoe and Waguespack (2011) show that attention to proposed Internet standards increases substantially when the presence of eminent author's name is revealed as opposed to hidden. Azoulay, Stuart and Wang (2012) show that citations increase to a researcher's prior body of work after the researcher becomes a Howard Hughes Medical Investigator, a high-status award in

⁴ For example, the increasing age at which biomedical researchers achieve their first NIH grant is well known, and may follow from the rising 'burden of knowledge' and teamwork that is increasingly ubiquitous in research and innovation (Jones 2010). Former NIH director Elias Zerhouni described the rising age at which researchers receive their first NIH grant as the most important challenge facing US science agencies (Kaiser 2008).

the biomedical sciences. Both studies indicate that positive reputational shocks can improve community awareness or perceptions of the scholar's existing output.

This paper departs from prior literature by emphasizing how reputation works in teams. The setting of team science allows us to examine not just how established reputations influence community viewpoints, but how differential reputations in the team influence individual-specific consequences. We thus embrace the centerpiece of Merton's seminal analysis, examining the potential entanglement of reputations, where eminent individuals may experience better consequences, but at the expense of others.

Our setting also appears original to our knowledge in emphasizing the consequences not of "good" events, but rather of team-produced catastrophes. Specifically, we consider consequences for researchers when a piece of their team-authored work is discovered to be false. The above discussion suggests several hypotheses about how prior reputations may influence reactions to these events. The communication hypothesis, normally an advantage, suggests that eminence may attract extra attention to the article retraction and thus amplify consequences for the authors involved. The credit hypothesis suggests two distinct alternatives. On the one hand, a strong reputation may protect an author in case of falsehood, where the community infers that a junior author was responsible for the problem. Thus the "rich get richer" aspect of the Matthew Effect may work in reverse, with eminence not only attracting good credit but also deflecting bad credit. On the other hand, the credit hypothesis may suggest that the community sees an eminent author as being "in charge" and directing events, in which case the eminent author may take the blame for mistakes, just as they get credit for successes. Other mechanisms may also bear on community reactions.⁵

Given a rich set of plausible mechanisms, we treat our analysis primarily as an empirical question and seek to establish first-order facts. Having presented these facts, we then return to theory more formally in Section V and provide a Bayesian

⁵ For example, team leaders may actively accept or deflect blame, and communities may follow norms in whether they choose to blame leaders. Across various organizational settings one can find examples of leaders who are fired for failures that occur under the "leader's watch", and contrasting examples where leaders scapegoat underlings.

interpretation that emphasizes the credit-inference aspects of the problem, where strong prior beliefs can insulate one's own reputation and deflect consequence onto others.

Whether or not our results provide guidance to many other team production settings, science is an important setting in its own right. Knowledge production, a foundation of economic growth, is increasingly done in teams across virtually all fields of science and engineering, social sciences, and patenting, and team-authored papers are increasingly likely to be the source of high impact work (Wuchty et al. 2007, Jones 2009). Thus, the classic ideas of Merton's Matthew Effect, should they be operating, are of increasing relevance to understanding the progress of science. Separately, article retractions are increasingly common and of growing concern among scientists and science institutions, including journals and funding agencies (Furman, Jensen and Murray 2012; Fang, Steen, and Casadevall 2012; Azoulay, Furman, Krieger, and Murray 2012; Lu et al. 2013).

The setting of science also offers useful empirical features for operationalizing reputational concepts and community responses. Modern databases of research articles provide codified outputs (papers) and codified measures of community use (citations) that allow rich opportunities to examine these classic ideas. We turn now to the data, empirical design, and results.

3. Data and Empirical Framework

Our data comes from the largest known repository of scientific knowledge, the Web of Science (WOS) from Thomson Reuters, which now includes more than 25 million publications published in over 15,000 journals worldwide, beginning in 1945. This database includes detailed bibliographic information for each paper (authors, journal, publication year, etc.) and further defines the citation linkages between each paper. The WOS further provides retraction notices that describe the time and reasons for each retraction and whether the errors are reported by the authors.

3.1. Treated Papers

In our study, we focus on changes in citations to an author's *prior published work*. We focus on prior work because this work is in a fixed published form, allowing us to isolate changes in usage of this work from changes in the work itself. Moreover, focusing on prior published work allows us to construct counterfactual cases by matching the prior work to other papers in the WOS that followed similar citation profiles prior to the retraction event. We refer to each prior publication by authors involved in the retraction as a treated paper.

Lu et. al (2013) show that retractions trigger citation losses to an author's prior work but also show that these penalties disappear if the author(s) self-report the error. Therefore, to examine how retraction affects authors by differential status, our sample focuses on retraction cases where scientific errors were not self-reported. For simplicity, we further restrict the sample to singular retraction cases, where an author is involved in only one retraction between 1993 and 2009.⁶ In this sample period we located 513 singular retraction events and 95% of these retracted papers (489) were written by more than one author. Among these team-authored retractions, 57.3% (280) were not self-reported, 32.3% (158) were self-reported, and 10.4% (51) had unclear or unknown retraction reasons. For each of the team-authored cases where retraction was not self-reported, we identified the authors' prior work published before the retraction.⁷ Changes in citations to these papers are the objects of our empirical analysis.

3.2 Control Papers

Because citation patterns differ across disciplines and by time since publication, we construct a control group to match each "treated" paper in the pre-retraction period. The underlying assumption is that both treated and control articles will continue the same course of citation patterns if there were no retraction influencing the treated

⁶ That is, we do not consider the (more extreme) cases where an author is revealed to have produced many false works, often entire bodies of works. These events are interesting but distinct in terms of the magnitude of the reputational consequences, the certainty about the guilty party, and the potentially diffuse timing of the shock(s), which makes such cases less relevant to the questions of interest and less amenable to the regression framework we employ.

⁷ The procedure we use to determine an author's prior work is described in Appendix A.

paper. This methodology draws on an identification approach first used in the context of scientific outputs by Furman and Stern (2011).

For a treated paper i published in field f and year p , we search for control papers within the same field and the same publication year. Using the WOS, we are able to search across millions of papers to find controls that are minimally distant within the same field, where field is defined by the 252 WOS field categories. In particular, for each non-treated paper j in this pool, we define the arithmetic distance between i and j as

$$AD_{ij} = \sum_{t=p}^{r-1} (c_{it} - c_{jt}) \quad (2)$$

and the Euclidean distance between i and j as:

$$ED_{ij} = \left[\sum_{t=p}^{r-1} (c_{it} - c_{jt})^2 \right]^{1/2} \quad (3)$$

where c_{it} denotes the citations paper i receives in year t and r is the year of retraction. Both distances attempt to measure the citation discrepancy between paper i and paper j , but arithmetic distance AD_{ij} allows for positive and negative differences to offset each other while Euclidean distance ED_{ij} is direction-free.

The quality of control group matching is assessed in Figure A1. Because we access the entire WOS, we can find substantially closer controls than is normally the case in other empirical applications of this treatment-control methodology (Furman and Stern 2011; Furman, Jensen and Murray 2012; Azoulay, Furman, Krieger, Murray 2012). For example, focusing on the ten papers with the lowest Euclidean distance to a treated paper, the upper-left panel of Figure A1 shows that the average Euclidean distance between the ten controls and the treated paper has high density around zero. The density drops smoothly at higher distances except for the bin of 50 or more (which is driven by some retracted papers that were exceptionally highly cited before retraction). As shown in the bottom-left panel of Figure A1, the average arithmetic distance between these ten controls and the treated paper has substantially more density on the negative side, so that these controls on average underestimate the citation flow of the treated papers. Focusing instead on the single control paper with the lowest Euclidean

distance, we are able to find a perfect match for 36.1% of the treated papers. When we cannot find a perfect match, the arithmetic distance of the single best control is negative on average, though it is more evenly distributed on both sides of zero than the ten-control sample.

To achieve a sample that balances close matches with sample size, we consider the two nearest neighbors, one from above (with positive AD) and one from below (with negative AD). As shown in the bottom-right panel of Figure A1, the density of the average arithmetic distance of these two controls is either exactly zero or concentrated in the neighborhood of zero. In particular, the two nearest neighbors now yield an average of zero arithmetic distance for a large share (68.5%) of our treated papers. This sample, with zero distance, is the main sample used in our analysis. In practice, we have 276 retraction events where authors have closely-matched prior work.⁸

Overall, by focusing on these 276 team-authored, single retraction events that were not self-reported, our sample includes 732 authors. The mean number of prior publications for these authors is 24.5. The mean number of prior publications for these authors where the two nearest-neighbor controls have zero average arithmetic distance is 16.8 giving a main treatment sample of 12,290 prior publications. Focusing on this sample, with each treatment paper and its two controls, the estimation sample includes 419,230 paper-year observations. Note that some prior publications will be counted more than once if multiple authors in the sample collaborated on them.⁹

3.3 Definitions of Author Status

⁸ Recall that there are 280 retraction cases of team-authored, single retractions where the authors do not self-report the error, thus we lose four events by focusing on prior publications that have close control matches prior the retraction event.

⁹ In practice, the estimation sample of 12,290 prior publications from retraction authors is constituted by 10,209 unique prior publications, some of which are shared by multiple retraction authors. We cluster standard errors by the retraction event (i.e. the 276 cases) to allow for correlated shocks across the prior work within a given author and across authors involved in the same retraction event.

We construct three standard measures for an author's status: publication counts, total citations received, and the h-index.¹⁰ The h-index (Hirsch 2005) attempts to account for publication quantity and quality in a single measure and is defined as follows: the number h is the largest scalar for a given scholar such that the scholar has published h papers each of which has been cited at least h times. These measures, which are commonly used as indications of eminence in the scientific community, are calculated using the papers and citations within the WOS. They are calculated for each author in the year just prior to the retraction event, based on their publication record up to that time.

Taking each treated author as an observation, Figure A2 plots the distribution of the h-index at the time of retraction. Consistent with the previous literature, the distribution is positively skewed, with a long tail on high status (MacRoberts and MacRoberts 1989, Selgen 1992). Similar skewness exists for paper counts and total citations. In the main part of our statistical analysis, we define the "absolute eminence" of an author using the continuous measures of paper counts, total citations, or h-index. As alternative measures, we also define simpler dummy variables to indicate whether an author is in the top 5th or top 10th percentile of a status distribution.

Because we focus on retractions of team-authored papers, we also define relative measures of social status based on whether an author has the highest or second highest social status in the team at the time of retraction. These authors are referred to as "relatively eminent." Compared to the absolute measure of author status, relative eminence helps us examine differential status within a team, even if all team members have high status or low status in absolute terms. The relative eminence measure can also help filter out the heterogeneity of social status measures across different academic fields.

3.4 Summary Statistics

¹⁰ In the regressions, we measure total prior publications in units of 1,000, total prior citations in units of 10,000 and prior h-index in units of 100.

Table 1 provides two panels of summary statistics: the first panel, at the author level, considers the status of each treated author at the time of retraction; the second panel, at the paper level, considers summary statistics for the retracted papers and prior work. The distribution of author measures (Panel A) shows that authors of a retracted paper had, at the time of retraction, a mean of 24 prior publications, 1,071 citations, and an h-index of 10. Whether measured by total counts of prior work, total counts of citation, or h-index, these author status measures appear dispersed and right-skewed.

Among the prior publications of these authors (Panel B), 45.5% were published in the 2000s, 40.0% were published in the 1990s, and 14.5% were published in the 1980s. The mean yearly citation count for the prior publications is 2.5. With our sample ending in 2009, the mean age of a prior publication in 2009 is 11.6 years. The mean age of a prior publication from an author in the year that author experiences a retraction is 8.5 years.¹¹

3.5 Estimation Equation

Our identification strategy employs differences-in-differences. We examine the citation effects of retraction shocks comparing the pre-post differences for treatment papers with the pre-post differences for control papers, while further comparing these differences across authors with different status. The regression model is

$$\Pr(y_{iat}) = f(\alpha_{ia} + \mu_t + \beta_1 \cdot Treat_i \cdot Post_{kt} + \beta_2 \cdot Status_a \cdot Treat_i \cdot Post_{kt} + \beta_3 \cdot Status_a \cdot Post_{kt} + \beta_4 \cdot Post_{kt}) \quad (4)$$

where i indexes article, a indexes author, t indexes year since publication, and k indicates a treatment-control paper group. The dependent variable, y , denotes counts of citations to article i at time t for author a . Fixed effects for each paper and author with retraction (α_{ia}) and each year since publication (μ_t) capture the mean citation pattern of articles. $Treat_i$ is a dummy variable that equals 1 if article i is a treatment paper, and

¹¹ With the rapid increase in retraction rates over the last decade (Fang et al. 2012, Lu et al. 2012), most retraction events provide a relatively brief window ex-post to observe ongoing citation behavior; thus, the regression analysis is primarily driven by citation responses to retraction events in the initial few years. We will explore effects on both recent and older publications below.

$Post_{kt}$ is a dummy variable that equals 1 if year t is after the retraction event for a given treatment and control group k . $Status_a$ measures the status of the treated author measured at the year just prior to the retraction.¹²

The coefficient β_1 captures the effect of the retraction shock on citations to prior work of non- eminent authors, compared to closely-matched control papers. The coefficient β_2 captures any difference in the effect for eminent authors compared to the non- eminent coauthors. We estimate (4) using the standard Poisson model for count data. While there are 10,209 unique prior publications in the treated sample, to be conservative we cluster the standard errors by the retraction event, giving 276 paper groups.¹³

The key identification assumption is that the prior work would continue the same course of citations as its control papers had the retraction not occurred. To the extent that this assumption may be less valid if the prior work is published close to the retraction time and therefore provides a shorter time window for matching control papers, we can exclude such cases as a robustness check and test whether the results change.

4. Results

As a first look at the citation patterns, Figure 1 shows the citation flows to prior publications before and after retraction, separating the data by author status. On the horizontal axis, zero demarcates the year of retraction. The solid blue line shows treated papers, and the dashed red line shows control papers. In the upper panel we consider absolute reputation, separating out those authors in the top 10th percentile of the h-index among all treated authors. The bottom panel repeats this exercise using relative reputation, distinguishing the highest-status author among the authors of the

¹² Note that the interaction term $Status_a * Treat_i$ is absorbed by the paper-author fixed effect (α_{ia}).

¹³ This approach allows arbitrary correlations in the errors across time for a given treated paper, across treated papers by the same author, and across all treated papers by distinct authors who were later involved in the same retraction event. A less conservative approach clusters papers based on the prior publication treatment-control group. Statistical precision with this latter approach is, not surprisingly, greater; these results are discussed briefly in Section 4.2.3 below.

retracted paper. These graphs suggest that the post-retraction citation decline is noticeably negative for ordinary authors, while eminent authors experience no citation loss. The rest of this section analyzes the data using regression models, presents our central findings, and considers various robustness checks.

4.1 Main Results on Author Status

Pooling the data in our sample across authors with different status, we first confirm that retraction has a significant negative spillover effect on citations to the authors' prior work. The regression results are presented in Figure 2. We see that, compared to the control papers, the annual flow of citations to prior publications falls 4.8% ($p < .0001$) in the first two years after the retraction and 13.0% ($p < 0.0001$) five or more years after the retraction. This suggests that retractions lead to substantial citation declines to prior work of team authors, which is consistent with the results shown in Lu et. al (2013).

4.1.1 Absolute Status

Table 2 reports results from our main specification. We highlight the differences-in-differences coefficient on *treated*post* retraction ($t \geq 1$) and the relative effect on individuals with higher status from the coefficient on *author status * treated * post* ($t \geq 1$).¹⁴ The latter indicates whether a treated author with greater absolute status at the time of retraction experiences different citation consequences for their prior work. There are three columns in the table, differing by measures of absolute status, using total prior publications, total prior citations, and the h-index respectively.

All measures show that the main effect (for those with low absolute status) is negative and statistically significant. Looking at the interaction with higher-status authors, the three continuous measures of reputation show that higher absolute status

¹⁴ We separate out the retraction year itself ($t=0$) because the exact time of retraction could occur early or late within the year.

offsets the negative main effect, with statistically significant interactions when using total prior citations or the h-index to measure status. Focusing on column (3), the coefficient indicates that a one standard deviation increase in prior h-index results in 4.9 percentage point smaller reduction in citations per year per paper due to retraction.¹⁵ This finding suggests that having higher status at the time of retraction may help alleviate the reputational harm due to retraction. Being eminent suggests a protective effect.

4.1.2 Status Relative to Coauthors

Beyond one's own absolute status, we further consider the implications of coauthors' relative status, as emphasized by Merton (1968). To do this, we conduct two sets of tests. In the first set of tests, we separate out those authors who have the highest status on the team, even if they don't have high status in an absolute sense. In particular, we define a dummy equal to one if a treated author has the highest status measure or, separately, if the author is among the two highest status individuals on the team. As before, status is measured in the year prior to the retraction and is alternatively examined using three different measures: the total number of prior publications, the total citations received, and the h-index.

Table 3 reports results, now measuring author status relative to other authors of the retracted teamwork. In columns (1)-(3) we separate out the highest-status author on the team. In columns (4)-(6) we separate out the two highest-status authors on the team.¹⁶

¹⁵ The standard deviation of the h-index at the paper level is 22.8 and the h-index regression variable is measured in units of 100. Thus a one standard deviation increase in the h-index translates to a $0.214 \times 22.8 / 100 = 4.9$ percentage point smaller citation loss per prior publication. Note that the standard deviation of the h-index across treated papers is larger than across treated authors (Table 1A). This follows because more prolific authors have more prior papers.

¹⁶ Recall that our sample includes only team-authored retracted papers. Among the retracted papers, 93% have three or more authors. To keep the sample identical across analyses, we continue to include the 7% of retracted papers with two-authors in columns (4)-(6). Limiting the sample to retracted papers with three or more papers produces virtually identical results in magnitude and statistical significance. Results are available upon request.

As before, the main effect for those with low relative status is negative and statistically significant across all specifications. When looking at the highest relative status author (Columns 1-3), we consistently see large, positive point estimates but the statistical significance is relatively weak at conventional levels.¹⁷ When looking at the two authors with highest relative status (Columns 4-6), we see larger point estimates and greater statistical significance across the measures. Moreover, the estimates for relatively low-status authors become increasingly negative, which suggests that looking at the top two individuals may divide high and low status individuals more precisely within the typical team.

In the second set of tests, we generalize the empirical model (4) to consider more textured team configurations. In particular, using binary absolute status measures (the top 10 percentile as the cutoff), we can consider the effects of retraction given four different status configurations among the authors of the retracted paper. These regressions include dummy variables to indicate whether (1) own status is ordinary and the highest-status coauthor is ordinary, (2) own status is ordinary but a coauthor is eminent, (3) own status is eminent and the highest-status coauthor is ordinary, and (4) own status and a coauthor are both eminent (the omitted category in the regression). Here, the coauthor refers to the best coauthor in a team. The results are presented in Table 4, Columns (1)-(3), with each column using a different measure of status: total publications, total citations, and the h-index. We see that the spillover effect on prior work is largest when one's own status is ordinary and one is in the presence of an eminent coauthor. This finding generalizes across the status measures with varying statistical significance. Taking column (3), for the h-index, the loss on prior work is 15.2% larger (i.e., $1 - \exp(-.165)$) when you are ordinary and your coauthor is eminent, compared to the baseline where you were also eminent yourself. Indeed, being eminent yourself suggests little citation losses to your prior work and regardless of the status of your coauthors, which is seen both in the main effect (you and a coauthor are eminent)

¹⁷ These results strengthen when looking at alternative specifications in Section 4.2.

and in the interaction effect where you are eminent and your highest status coauthor is not.

Taken together, the results in Table 3 and Table 4 show a consistent pattern. After retraction, ordinary authors experience large citation losses to their prior work, especially when working with an eminent coauthor. Eminent authors, by contrast, show little citation losses to their prior work, regardless of the status of their coauthors. A variety of additional tests, discussed below, tend to further support these results and tend to strengthen their magnitudes or statistical precision.

4.2 Additional Tests and Robustness Checks

We consider here several additional tests that explore the robustness of the above results and can further sharpen the empirical findings.

4.2.1 Self Citations

Retractions may also affect future publishing prospects, and differentially for eminent and non-eminent authors. The decline in citations to prior work might then potentially reflect less a direct community response and more a decline in the capacity of the authors to cite their own prior work, once any differential retraction effects on an author's career take hold. To further focus on the community response, we reconsider the analysis excluding self-citations from the citation counts. These results are presented in Table 5. The findings are very similar to Tables 2, 3 and 4, no matter whether we use the absolute or relative status measures. Interestingly, the magnitude in citation reduction for the main effects becomes slightly larger for both the absolute and relative measures. This finding of larger citation losses for lower-status authors, when netting out self-citations, further implies that the negative spillover effect on prior work comes from the broader community.

4.2.2 Old Papers

Older papers may receive fewer ongoing citations, and no paper can receive less than zero citations after retraction. Because eminent authors are more senior and may have an older distribution of papers than ordinary authors do, this tendency could contribute to the smaller citation reductions for eminent authors.

Figure A3 shows the citation trajectories for our treated papers. The average citations for treated papers fall to two in the tenth year since publication and fall to one in the fifteenth year since publication. Given these facts, we reconsider our analysis excluding prior articles published more than 10 years earlier than the retraction year. As a result, 68.9% of treated papers and 59.1% of paper-year observations are kept in the subsample.

As shown in Table 6, results estimated on this subsample remain robust. Citations losses remain much larger for low-status authors after retraction. For example, according to Column (6) citations fall by 14.4% (i.e., $1 - \exp(-.155)$) for low-status authors after retraction and the percentage difference between high- and low-status researchers is 9.9% ($1 - \exp(-.104)$). If the old paper hypothesis holds, the coefficient of $Treated * Post(t \geq 1)$ should be more negative and the differences between high- and low-reputation authors would be smaller after old papers are excluded from the citation counts. Both numbers shown in Column (6) of Table 6 are similar to the corresponding ones in Table 3, as is broadly the case when comparing the results in Table 6 with the earlier results in Tables 2, 3, and 4. These findings are inconsistent with the old paper hypothesis and to the contrary, like the analysis net of self-citations, such a sampling restriction appears to modestly strengthen the results.

4.2.3 Sample and Regression Model

We further conduct a series of robustness checks by estimating different samples and different models. First, we replace our Poisson estimation with OLS estimation. The OLS results are reported in Table A1 and appear broadly similar to the Poisson results. Second, we explore the main results again in Table A2 clustering the standard-errors instead by treatment-control paper group instead of retraction event, which is seen to

strengthen the statistical precision. Third, we further consider a restricted sample where all publications are being positively cited at the time of retraction. This issue is different from the old paper hypothesis because zero citations could occur soon after publication, especially for ordinary authors who do not have many high quality publications. To deal with this issue, we exclude all prior work that has zero citations in the year before retraction. As shown in Table A3, results remain robust in those still-cited papers. Fourth, we separate out prior work that has a short citation history before retraction, which could hurt our ability to find effective counterfactual controls. We address this issue by excluding all prior work published within three years before retraction. Results are shown in Table A4 and appear similar to but slightly stronger than our baseline specification.

Finally, we consider a Placebo exercise to see whether the evolution of control paper citations is sensitive to status in the absence of retraction. In particular, using our control papers, we examine whether papers matched according to very similar initial citation patterns also have similar later citation patterns regardless of status. We find that status does not predict future citation paths, conditional on initially similar citation paths, as detailed in Table A5. This analysis further suggests that our control strategy is effective for estimating counterfactual citation paths in the absence of retraction.

Overall, the results remain robust and these additional analyses further support our main finding: ordinary authors experience much greater consequences than eminent authors when a paper is retracted, as measured by the tendency for the community to continue to cite their prior work.

5. Interpretations and Discussion

The above empirical analyses establish several striking facts regarding the retraction shocks and their differential effects across team members. We call these results a “Reverse Matthew Effect”, as they echo the “rich get richer” idea of Merton’s classic Matthew Effect, only now in the reverse case where we consider bad events. We find that retraction shocks lead to substantial declines in citations to the prior work of

ordinary coauthors. By contrast, for eminent coauthors of the retracted publication, retraction shocks provoke little if any citation loss to their prior body of work. Furthermore, citation losses for ordinary coauthors are especially severe in the presence of an eminent coauthor on the retracted publication.

This section further discusses the empirical results in light of the classic mechanisms that Merton proposed. Returning to Merton’s credit mechanism, we first formalize the idea that the community makes ex-post inferences about individual contributions in team settings given prior reputations and the uncertainty over who was responsible for the output. A simple, Bayesian model of this mechanism is shown to provide a parsimonious, candidate explanation for the empirical results. We further discuss an alternative credit mechanism and Merton’s communication hypothesis in light of the empirical findings.

5.1 A Model

Let there be two types of agents, who differ in their tendency to produce “bad” output. The community does not observe an individual's type directly but rather makes inferences about it by observing the individual’s output. The community's belief about the individual's type characterizes that individual's reputation.

In particular, let an agent i produce bad output with probability $p_i \in \{p_L, p_H\}$, where $0 \leq p_L \leq p_H \leq 1$, so that type L individuals produce bad output with relatively low probability and type H individuals produce bad output with relatively high probability. Denote the community's belief about an individual's type as R_i where

$$R_i = \Pr[L_i]$$

That is, the person has a perceived ‘reputation’, which is represented by the probability the community assigns to that person being the low-error type, L. Finally, let a piece of output be denoted $y \in \{T, F\}$, where T indicates that the output has no known errors while F indicates that the output is “bad”. In our empirical setting, an “F” event indicates a retraction, and a high R_i indicates a well-established author, someone who has established a reputation for producing good rather than bad output.

In summary, the background probability of producing bad output, p_i , depends on the author's type (L or H). How to distinguish the type given the observed output is the heart of the inference problem.

5.1.1 Solo Production

To develop basic intuition, first consider the reputational updating for an individual who, working alone, produced a bad piece of output. Let the individual have a given prior reputation, R_i . Bayes rule says that the posterior belief about i's type, which we denote R_i' is

$$R_i' = \Pr[L_i|F] = \frac{\Pr[F|L_i] \Pr[L_i]}{\Pr[F]}$$

Using the law of total probability in the denominator and definitions above, we can thus express the reputational change upon retraction as

$$\frac{R_i'}{R_i} = \frac{p_L}{p_L R_i + p_H (1 - R_i)}$$

Recalling that $p_L \leq p_H$, it follows that a retraction can only worsen the individual's reputation ($R_i' \leq R_i$). It also follows that the percentage change in the individual's reputation is declining in R_i . In the extreme case, where $R_i = 1$, the individual is fully protected from the reputational consequences of retraction; as is standard with a Bayesian model, having a very tight prior about the individual means that new events will have little further effect on beliefs.

5.1.2 Team Production

We now consider the richer case of team production, which allows us to characterize a "Reverse Matthew Effect". In particular, let the piece of output be produced by a team of two people, indexed $i \in \{1,2\}$, who have independent probabilities of making a mistake and independent priors.¹⁸ As above, let the output

¹⁸ The assumption of independent priors is made for simplicity. In team production, individuals may have produced together before and thus the priors may not be fully independent. While that case may be interesting, our goal here is to provide the simplest characterization for our empirical results.

turn out to be “bad”. By Bayes’ Rule, the posterior belief about the quality of individual 1 can be written

$$R'_1 = \Pr[L_1|F] = \frac{\Pr[F|L_1, L_2] \Pr[L_1, L_2] + \Pr[F|L_1, H_2] \Pr[L_1, H_2]}{\Pr[F]}$$

In other words, we now need to integrate out over the possible cases for individual 2.

Using the law of total probability to determine $\Pr[F]$, the definitions above to determine the individual probability terms, and some algebra, we can write the change in reputation as

$$\frac{R'_1}{R_1} = \frac{1}{R_1 + (1 - R_1) \frac{bR_2 + a(1-R_2)}{cR_2 + b(1-R_2)}} \quad (5)$$

where $a = \Pr[F|H_1, H_2] = 1 - (1 - p_H)^2$, $b = \Pr[F|H_1, L_2] = 1 - (1 - p_L)(1 - p_H)$, $c = \Pr[F|L_1, L_2] = 1 - (1 - p_L)^2$, and we note that $1 \geq a \geq b \geq c \geq 0$.

This expression presents four central results, encapsulated in the following Lemma.

Lemma (i) $R'_1 \leq R_1$; (ii) $\frac{\partial(R'_1/R_1)}{\partial R_1} \geq 0$; (iii) $\frac{\partial(R'_1/R_1)}{\partial R_2} \leq 0$; and (iv) $\frac{\partial}{\partial R_1} \left(\frac{\partial(R'_1/R_1)}{\partial R_2} \right) \geq 0$.

The proof is given in the appendix.

These results capture the empirical findings and provide some precise intuition for them. The first result states that reputational losses from a retraction are negative. This result corresponds to the broad finding where authors experience citation losses on average to their existing work and the finding that no authors appear to actively benefit from a retraction. The second result states that a high reputation acts to limit the reputational decline from the retraction. This result corresponds to the findings in Table 2, where authors with a lower absolute reputation experience more negative consequences on average compared to eminent authors. This finding gives the first appearance of a “Reverse Matthew Effect”, where eminence appears protective in the context of negative events.

The last two results focus on the reputational entanglement across authors that may emerge in a teamwork setting. The third result states that the greater the reputation of your coauthor, the worse the effect on you. Thus, the Bayesian model

predicts that the presence of an eminent coauthor exacerbates the reputational losses for the other author. At the same time, the fourth result shows that eminence is protective against this spillover effect. Thus, while an eminent coauthor can hurt you, it hurts you less if you yourself are eminent. These theoretical results are closely consistent with the findings in Table 4, where ordinary authors experience worse effects the more eminent the coauthor, yet eminent authors see little effect from eminent coauthors.

These results are all intuitive in a Bayesian context, where the community is trying to infer the source of a mistake and must adjudicate between the authors and the background chance of a mistake. A well-established reputation deflects blame away from you and toward both your coauthor and background bad luck. If the coauthor also has a well-established reputation, then the community will tend to blame background bad luck, and both authors face relatively mild consequences. An unformed reputation, however, attracts blame, and the more so the better your coauthor's reputation. The credit inference problem that animates Merton's Matthew Effect (Merton 1968) in the context of team production can thus provide a natural and parsimonious interpretation of the results.

5.2. An Alternative Credit Inference Hypothesis

Within the class of credit inference explanations, an alternative inference problem involves task allocation within the team. In particular, one may argue that science teams feature a hierarchal nature; eminent authors typically lead in the conceptual design of the research rather than in the technical analysis, where problems are more likely to emerge. In this view, eminent authors may receive less blame when retraction occurs because they are seen as unlikely to be responsible for the relevant tasks.

One way to test this idea is to control for position in the author list for the retracted paper. Noting that positioning in the author list typically informs the hierarchy of the team in the science and engineering, we reconsider our main results adding dummies variables for the last author (usually the principle investigator and/or

laboratory head) and middle authors (who play lesser roles). As shown in Table 7, adding such author-position fixed effects to the regression model has little effect on the results; these author position fixed effects are highly insignificant, while the coefficients on the status measures remain similar in magnitude and significance as before.

Another way to test this idea is to examine citation effects based not on author status at the time of the retraction but at the time the research was conducted, when task allocation would be determined. To do so, we constructed past-status measures using the status of an author in the year the problem paper was published. Then we examined both types of author status (at the time of retraction and at the time of publication) in the regression. For ease of interpretation, both types of status are measured by a dummy for whether the absolute status is in the top 10 percentile of all treated authors at that time. As shown in the first three columns of Table 8, being eminent at time of retraction substantially reduces the citation losses using two of the three status measures, while being eminent at time of publication does not. This result appears inconsistent with a task allocation hypothesis. The last three columns of Table 7 restrict the sample to authors who had ordinary status when the problem paper was published. Some of these authors became high-status and others remained ordinary by the time of retraction. Results shown in the last three columns of Table 7 suggest that ordinary authors who became eminent later, measured by total publications or h-index, see little if any citation loss. These results further suggest that task allocation does not appear to be a key explanation for our main findings.

5.3 The Communication Hypothesis

Merton's Matthew Effect also emphasizes a "communication" hypothesis, where eminence attracts attention to the output, for which there is evidence in the literature (Simcoe and Waguespack 2011, Azoulay, Stuart, and Wang 2012). In the standard Matthew Effect, which considers "good" events, this communication effect may help the less established author, offsetting the credit sharing issue. Namely, even if the less established author receives little credit *share*, a widely noticed output can make this

little share larger in absolute terms. With a “bad” event, the communication hypothesis could, by contrast, makes things worse for affected authors, as the presence of an eminent author may make bad events more widely noticed.

While a communication mechanism may be operating in our context, it does not appear capable of providing an alternative explanation for the results. Namely, were this mechanism all that was happening, then eminence should worsen the citation losses in general. Given that we find the opposite result -- that ordinary authors experience substantially worse effects than eminent authors -- the communication hypothesis does not appear to dominate. Nonetheless, the basic communication mechanism may still be operating in tandem with other forces. For example, if high status is protective from a Bayesian perspective, and low status is not (as in Section 5.1) then the communication channel may worsen things more for the less eminent in the presence of eminent coauthors, exacerbating the credit inference effects. It is also possible that, in our empirical setting, retractions are sufficiently well noticed that the marginal additional communication effect of eminence is small. In that sense, catastrophes may be settings where credit inference mechanisms dominate communication mechanisms; for “good” events, the balance of these forces may be different.

6. Conclusion

We have considered a natural experiment to assess the consequences of retraction. Our results demonstrate asymmetry: Eminent authors show little or no change in citations to their prior work after a coauthored retraction, while less eminent coauthors experience large citation losses, and especially in the presence of an eminent coauthor. These and other results are consistent with a simple Bayesian model that operationalizes the “credit inference” problem in teamwork settings that animates Merton’s canonical Matthew Effect. Not only do the rich get richer, when riches are to be had, but the poor get poorer when catastrophe strikes.

Team production now comprises the vast majority of papers in the sciences and engineering. Therefore, issues of credit sharing become more acute. Especially for junior scientists, who increasingly establish their individual reputations exclusively through team-authored outputs, the Matthew Effect presents a difficult challenge. If established authors can both take credit for successes and avoid discredit from failures, the junior author may take substantially longer to develop their own reputation while facing greater career risks along the way. These features may act as entry barriers to scientific careers. More subtly, these concerns may influence how scientists choose collaborators, so that credit considerations turn scientists away from potentially productive teams. Junior researchers have to evaluate the tradeoff between the credit sharing effect and the positive effect that an eminent coauthor can bring in attention and citations. These issues are important areas for future work.

While our setting is scientific teamwork, the primitives of our setting – collaboration across individuals, uncertainty over output quality, and differential reputations, generalize across many production contexts. Damaging or catastrophic events in collaborative settings range from food poisoning and airplane crashes to surgical mishaps and accounting fraud. The science context, with its codified outputs (papers) and codified measures of community use (citations), provides one inroad, and a classically motivated one, to this more general phenomenon. Empirical and theoretical investigations that can improve our understanding of underlying mechanisms and their implications, in knowledge production and in many other production contexts, provide exciting areas for further study.

References

- Azoulay, Pierre; Toby Stuart; and Yanbo Wang (2012) “Matthew: Effect or Fable” NBER working paper series #18625.
- Azoulay, Pierre; Jeffrey L. Furman; Joshua L. Krieger, and Fiona E. Murray (2012): “Retractions” NBER working paper series #18499.
- Bacon, Francis (1620). *Novum Organum*.

- Bajari, Patrick and Ali Hortacsu (2004) "Economic Insights from Internet Auctions." *The Journal of Economic Literature*, 42(2): 457-486.
- Becker, Gary S. and Kevin M. Murphy (1992). "The Division of Labor, Coordination Costs, and Knowledge," *Quarterly Journal of Economics*, 107 (4), 1137-1160.
- Cabral, Luis and Ali Hortacsu (2010): "The Dynamics of Seller Reputation: Evidence from eBay" *The Journal of Industrial Economics* 58(1): 54-78.
- Cooper, David J. and John H. Kagel (2005): "Are Two Heads Better than One? Team versus Individual Play in Signaling Games" *The American Economic Review*, 95(3): 477-509.
- Dranove, David; Subramaniam Ramanarayanan and Yasutora Watanabe (2012): "Delivering Bad News: Market Responses to Negligence", *Journal of Law and Economics*, 55(1).
- Fang, Ferric, R. Grant Steen, and Arturo Casadevall (2012): "Misconduct accounts for the majority of retracted scientific publications" *PNAS*.
- Freeman, Richard, Ina Ganguli and Raviv Murciano-Goroff (2013). "Why and Wherefore of Increased Scientific Collaboration," mimeo, Harvard University.
- Furman, Jeffrey L. and Scott Stern (2011): "Climbing atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research" *American Economic Review* 101: 1933-1963.
- Furman, Jeffrey L., K. Jensen, and Fiona Murray (2012). "Governing Knowledge in the Scientific Community: Exploring the Role of Retractions in Biomedicine," *Research Policy* 41 (2): 276-290.
- Hamilton, Barton; Jack A. Nickerson and Hideo Owan (2003) "Empirical Analysis of the Impact of Teams on Productivity and Participation" *Journal of Political Economy* 111(3): 465-497.
- Hirsch, J. E. (2005). An Index to Quantify an Individual's Scientific Research Output," *Proceedings of the National Academy of Sciences* 102 (46) 16569-16572.
- Holmstrom, Bengt (1982). "Moral Hazard in Teams," *Bell Journal of Economics* 324-340.
- Jin, Ginger Z. and Phillip Leslie (2009): "Reputation Incentives for Restaurant Hygiene"

- American Economic Journal: Microeconomics* 1(1): 237-67.
- Johnson, Erin (2012): "Ability, Learning and the Career Path of Cardiac Specialists" MIT Sloan working paper.
- Jones, Benjamin F. (2009): "The Burden of Knowledge and the Death of the Renaissance Man: Is Innovation Getting Harder?" *Review of Economic Studies*. 76(1).
- Jones, Benjamin F. (2010). "As Science Evolves, How Can Science Policy?" NBER Innovation Policy and the Economy 11.
- Jones, Benjamin, Stefan Wuchty, and Brian Uzzi (2008). "Multi-university Research Teams: Shifting Impact, Geography, and Stratification in Science," *Science* 322, 1259-1262.
- Kaiser, Jocelyn (2008). "Zerhouni's Parting Message: Make Room for Young Scientists," *Science* 322, 834-35.
- Klein, Benjamin and Keith B. Leffler (1981): "The Role of Market Forces in Assuring Contractual Performance" *Journal of Political Economy* 89(4): 615-641.
- Lu, Susan Feng; Ginger Jin; Benjamin Jones; Brian Uzzi "The Prior Publication Penalty: An Investigation of Retraction and False Science", *working paper*, 2012.
- MacRoberts, M.H. And B.R. MacRoberts (1989) "Problems of Citation Analysis: A Critical Review", *Journal of the American Society for information Science*, 40(5): 342-349.
- Mas, Alexandre and Enrico Moretti (2009): "Peers at Work" *The American Economic Review*, 99(1): 112-145.
- Merton, Robert (1968) "The Matthew Effect in Science", *Science* 159 (3810): 56-63.
- Merton, Robert (1988) "The Matthew Effect in Science, II: Cumulative Advantage and the Symbolism of Intellectual Property" *ISIS* 79: 606-623.
- Pope, Devin (2009): "Reacting to Rankings: Evidence from 'America's Best Hospitals'" *Journal of Health Economics* 28(6), 1154-1165.
- Selgen, Per O. (1992) "The Skewness of Science", *Journal of the American Society for Information Science* 43(9): 628-638.
- Shapiro, Carl (1983): "Premiums for High Quality Products as Returns to Reputations"

- The Quarterly Journal of Economics* 98(4): 659-679.
- Simcoe, Tim and Dave Waguespack (2011). "Status, Quality, and Attention: What's in a (Missing) Name?" *Management Science* 57, 274-290.
- Smith, Adam (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*.
- Uzzi, Brian and Jarrett Spiro (2005). "Collaboration and Creativity: The Small World Problem," *American Journal of Sociology*, 111, 447-504.
- Uzzi, Brian, Satyam Mukherjee, Michael Stringer, and Ben Jones (2013). "Atypical Combinations and Scientific Impact," mimeo, Northwestern University.
- Wuchty, Stefan, Benjamin F Jones and Brian Uzzi (2007): "The Increasing Dominance of Teams in the Production of Knowledge" *Science*. 316(5827): 1036-1039.

Appendix A: Prior Work

We built the sample of prior work using the Web of Science database. Because different authors may share the same name, relying on the name alone to identify an author's body of work would result in an inaccurate sample. We therefore applied the following procedures, harnessing the citation network, to identify the authors' prior work.

- We compiled a list of retracted articles and obtained the names of authors for each article.
- We then exploited the citation network in the Web of Science to identify the articles cited by these authors, where the cited article shares the citing author's name.
 - Specifically, we start by tracing citations from each retracted article to prior articles by the same author, and then use the citations from these prior articles to other prior articles by the same author and so on up to a point when additional prior work is no longer available.
 - Next, we use the obtained prior work to trace forward this citation network and locate papers by the same author that cite these past publications.
 - We use the retraction year as a cutoff to identify the authors' work published before the retraction.
 - Note that we exclude any prior work that was retracted itself.
 - Some prior publications will be counted more than once if multiple authors in the sample collaborated on them.

The prior publications identified in this way are written by the same author and they should capture most of the prior works that this author has written on a topic related to the retracted work. It may fail to capture the papers that are written by the same person but in completely unrelated areas. Possibly, it will include authors that are distinct people but share the same name and work in the same research area, as defined

by the citation network, although simple estimations suggest that such mismatches are extremely unlikely.

Appendix B: Proof of Lemma

The Lemma is repeated here for convenience, with the proof following.

Lemma (i) $R'_1 \leq R_1$; (ii) $\frac{\partial(R'_1/R_1)}{\partial R_1} \geq 0$; (iii) $\frac{\partial(R'_1/R_1)}{\partial R_2} \leq 0$; and (iv) $\frac{\partial}{\partial R_1} \left(\frac{\partial(R'_1/R_1)}{\partial R_2} \right) \geq 0$.

Proof

Recall equation (5), which we write here as $R'_1/R_1 = \left[R_1 + (1 - R_1) \frac{bR_2 + a(1-R_2)}{cR_2 + b(1-R_2)} \right]^{-1}$.

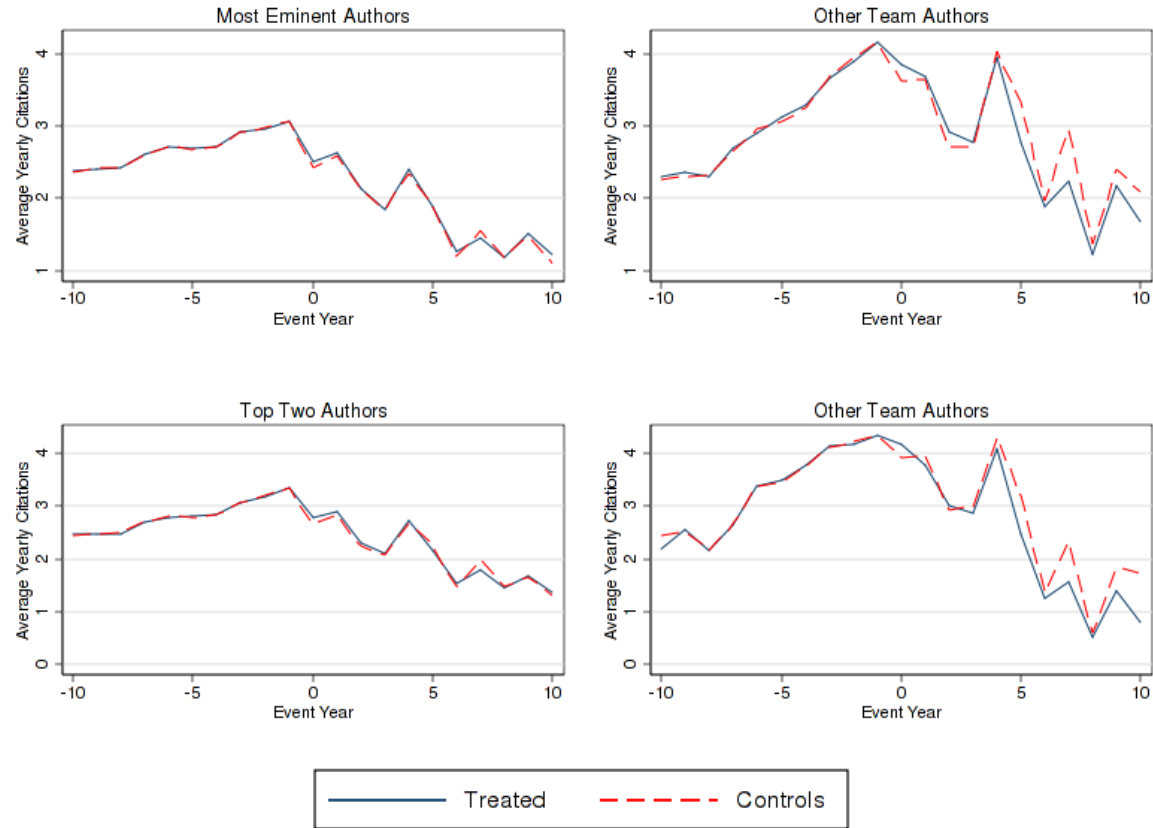
Result (i) follows by noting that $\frac{bR_2 + a(1-R_2)}{cR_2 + b(1-R_2)} \geq 1$. This ratio exceeds 1, by inspection, because $b \geq c$ and $a \geq b$.

Result (ii) also follows by inspection, noting again that $\frac{bR_2 + a(1-R_2)}{cR_2 + b(1-R_2)} \geq 1$.

Result (iii) follows if $\frac{\partial}{\partial R_2} \left(\frac{bR_2 + a(1-R_2)}{cR_2 + b(1-R_2)} \right) \geq 0$. It can be shown that $\frac{\partial}{\partial R_2} \left(\frac{bR_2 + a(1-R_2)}{cR_2 + b(1-R_2)} \right) = \frac{b^2 - ca}{(b + (c-b)R_2)^2}$, so that the sign of this derivative is the sign of $b^2 - ca$. Returning to the underlying definitions of a , b , and c (see main text) and writing $e = 1 - p_H$ and $f = 1 - p_L$, one can write $b^2 - ca = (e - f)^2 \geq 0$, proving the result.

Result (iv) follows by inspection of (5), given result (iii).

Figure 1: Citation before and after retraction, by author status



Note: The solid blue line indicates the treated papers (prior publications of authors involved in the retraction), and the dashed red line indicates control papers. In the top row, “Other Team Authors” are all but the most eminent author in the team of the retracted paper. In the bottom row, “Other Team Authors” are all but the two most eminent authors in the team of the retracted paper.

Figure 2: The effect of retraction on citations to an author's prior publications, compared to control papers, by year since the retraction event




Table 1: Summary statistics

Panel A: Unit of observation = author, treated only

Absolute Measures of Status	Definition	Obs	MEAN	SD	Min	Max
Prior Publications	total prior papers	732	24	46	1	452
Prior Citations	total prior citations	732	1071	3570	0	67946
Prior h-index	prior h-index	732	10	14	0	132

Panel B: Unit of observation = paper, treated only

	Retracted Papers	Prior Work
Paper Counts	276	10,209
% Published in 2000s	86.2%	45.5%
% Published in 1990s	13.8%	40.0%
% Published in 1980s	0%	14.5%
Yearly Mean Citation Count	3.3	2.5
Mean Age Since Publication ^(a)	5.3	11.6
Mean Age at Retraction ^(b)	2.2	8.5

(a) Age since publication is the difference between 2009 (the end of our sample) and the publication year; (b) Age at retraction is the difference between the year of the retraction event and the publication year.

Table 2: Effect of retraction on citations to prior work, by absolute status measures of the treated author at the time of retraction

Absolute Status of the treated author	Status Measures		
	Total # of prior papers	Total # of prior citations	prior H-index
	(1)	(2)	(3)
Treated*Post(t>=1)	-0.102** (0.044)	-0.109*** (0.036)	-0.135*** (0.048)
Author Status*Treated*Post(t>=1)	0.349 (0.313)	0.083** (0.035)	0.214** (0.109)
Treated*Post(t=0)	0.001 (0.025)	-0.003 (0.024)	-0.012 (0.031)
Author Status*Treated*Post(t=0)	0.036 (0.145)	0.014 (0.013)	0.052 (0.054)
Post(t>=1)	-0.112** (0.056)	-0.077 (0.060)	-0.058 (0.086)
Post(t=0)	-0.180*** (0.050)	-0.161*** (0.053)	-0.156** (0.063)
Author Status*Post(t>=1)	-0.565 (0.717)	-0.155 (0.106)	-0.344 (0.337)
Author Status*Post(t=0)	-0.770* (0.430)	-0.181** (0.076)	-0.310* (0.182)
Author-Paper Fixed Effects	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y
Observations	419,239	419,239	419,239
Number of papers	34,562	34,562	34,562

Author status refers to the absolute status of a treated author in the year prior to retraction. For interpreting regression coefficients, the total number of prior papers is measured in 1,000s; the total number of prior citations in measured in 10,000s and the h-index is measured in 100s. All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table 3: Effect of retraction on citations to prior work, by status of the treated author relative to his/her coauthors at the time of retraction

Status of a treated author relative to the other coauthors within the team	Discrete measures					
	Top 1 in Total # of prior work	Top1 in Total # of prior citations	Top 1 in h-index	Top 2 in Total # of prior work	Top2 in Total # of prior citations	Top2 in h-index
	(1)	(2)	(3)	(4)	(5)	(6)
Treated*Post(t>=1)	-0.114** (0.044)	-0.119*** (0.045)	-0.119*** (0.045)	-0.175*** (0.046)	-0.151*** (0.055)	-0.154*** (0.052)
Author Status*Treated*Post(t>=1)	0.065 (0.042)	0.074* (0.043)	0.072* (0.043)	0.121*** (0.046)	0.095* (0.056)	0.097* (0.053)
Treated*Post(t=0)	-0.004 (0.031)	-0.012 (0.032)	-0.009 (0.033)	-0.009 (0.038)	-0.018 (0.037)	-0.006 (0.037)
Author Status*Treated*Post(t=0)	0.014 (0.033)	0.026 (0.034)	0.020 (0.034)	0.016 (0.036)	0.026 (0.035)	0.012 (0.035)
Post(t>=1)	-0.125** (0.054)	-0.129** (0.055)	-0.120** (0.054)	-0.099* (0.057)	-0.116** (0.055)	-0.120** (0.056)
Post(t=0)	-0.263*** (0.066)	-0.263*** (0.065)	-0.269*** (0.067)	-0.241*** (0.078)	-0.263*** (0.069)	-0.269*** (0.071)
Author Status*Post(t>=1)	-0.069 (0.044)	-0.064 (0.046)	-0.077* (0.045)	-0.083* (0.050)	-0.064 (0.050)	-0.059 (0.050)
Author Status*Post(t=0)	0.020 (0.043)	0.021 (0.042)	0.028 (0.045)	-0.011 (0.055)	0.016 (0.046)	0.023 (0.048)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y
Observations	419,239	419,239	419,239	419,239	419,239	419,239
Number of papers	34,562	34,562	34,562	34,562	34,562	34,562

Author status refers to the dummy of whether a treated author had the highest status (“Top 1”) within the team or is among the two individuals with highest status (“Top 2”) in the year prior to retraction. All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table 4: Effect of retraction on citations to prior work, by own and coauthor reputation

Status configurations of own and co-authors in the retracted teamwork	All Authors		
	Total # of prior work	Total # of prior citations	Prior h-index
	(1)	(2)	(3)
Treated*Post(t>=1)	-0.016 (0.037)	-0.059 (0.076)	0.009 (0.029)
Self is eminent and Co-author is ordinary *Treated*Post(t>=1)	-0.029 (0.061)	-0.002 (0.093)	-0.056 (0.060)
Self is ordinary and Co-author is eminent *Treated*Post(t>=1)	-0.123* (0.067)	-0.126 (0.097)	-0.165** (0.082)
Self is ordinary and Co-author is ordinary *Treated*Post(t>=1)	-0.063 (0.064)	0.009 (0.089)	-0.101* (0.057)
Author-Paper Fixed Effects	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y
Observations	419,239	419,239	419,239
Number of papers	34,562	34,562	34,562

We classified the authors into four groups using dummy variables indicating whether (1) own status is ordinary and the highest-status coauthor is ordinary, (2) own status is ordinary but a coauthor is eminent, (3) own status is eminent and the highest-status coauthor is ordinary, and (4) own status and a coauthor are both eminent (the omitted category in the regression). Author status is measured in the year prior to retraction. All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. All regressions include all one-way and two-way interactions terms; we do not report those coefficients for brevity. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table 5: Effect of retraction on citations to prior work, excluding self-citations

Measure of Author Status	Absolute Status			Relative Status			Team Configuration		
	Total # of prior work	Total # of prior citations	h-index	Top 2 in Total # of prior work	Top2 in Total # of prior citations	Top2 in h-index	Total # of prior work	Total # of prior citations	h-index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated*Post(t>=1)	-0.127***	-0.148***	-0.175***	-0.205***	-0.186***	-0.186***	-0.059	-0.087	-0.031
	(0.046)	(0.039)	(0.051)	(0.049)	(0.058)	(0.055)	(0.060)	(0.078)	(0.046)
Author Status*Treated*Post(t>=1)	0.320	0.105***	0.255**	0.124**	0.103*	0.102*			
	(0.342)	(0.037)	(0.117)	(0.049)	(0.058)	(0.055)			
Self is eminent and Co-author is ordinary							-0.016	0.016	-0.027
*Treated*Post(t>=1)							(0.080)	(0.096)	(0.072)
Self is ordinary and Co-author is eminent							-0.135*	-0.147	-0.174*
*Treated*Post(t>=1)							(0.079)	(0.098)	(0.090)
Self is ordinary and Co-author is ordinary							-0.030	0.001	-0.092
*Treated*Post(t>=1)							(0.081)	(0.092)	(0.068)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	418,128	418,128	418,128	418,128	418,128	418,128	418,128	418,128	418,128
Number of papers	34,361	34,361	34,361	34,361	34,361	34,361	34,361	34,361	34,361

For interpreting regression coefficients in Columns (1)-(3), the total number of prior papers is measured in 1,000s; the total number of prior citations in measured in 10,000s and the h-index is measured in 100s. When author reputation is measured in relative terms, please look at the notes in Table 3 and 4. All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. All regressions include all one-way and two-way interactions terms; we do not report those coefficients for brevity. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table 6: Effect of retraction on citations to prior work, excluding old papers

Measure of Author Status	Absolute Status			Relative Status			Team Configuration		
	Total # of prior work	Total # of prior citations	h-index	Top 2 in Total # of prior work	Top2 in Total # of prior citations	Top2 in h-index	Total # of prior work	Total # of prior citations	h-index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated*Post(t>=1)	-0.099**	-0.105***	-0.129***	-0.178***	-0.153***	-0.155***	0.008	-0.052	0.016
	(0.046)	(0.038)	(0.050)	(0.046)	(0.053)	(0.051)	(0.033)	(0.090)	(0.032)
Author Status*Treated*Post(t>=1)	0.365	0.083**	0.214**	0.131***	0.102*	0.104**			
	(0.323)	(0.034)	(0.106)	(0.045)	(0.053)	(0.050)			
Self is eminent and Co-author is ordinary *Treated*Post(t>=1)							-0.047	-0.009	-0.056
							(0.057)	(0.102)	(0.059)
Self is ordinary and Co-author is eminent *Treated*Post(t>=1)							-0.143**	-0.129	-0.164**
							(0.066)	(0.104)	(0.083)
Self is ordinary and Co-author is ordinary *Treated*Post(t>=1)							-0.089	0.007	-0.107*
							(0.065)	(0.102)	(0.062)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	211,788	211,788	211,788	211,788	211,788	211,788	211,788	211,788	211,788
Number of papers	24,121	24,121	24,121	24,121	24,121	24,121	24,121	24,121	24,121

For interpreting regression coefficients in Columns (1)-(3), the total number of prior papers is measured in 1,000s; the total number of prior citations in measured in 10,000s and the h-index is measured in 100s. When author reputation is measured in relative terms, please look at the notes in Table 3 and 4. All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. All regressions include all one-way and two-way interactions terms; we do not report those coefficients for brevity. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table 7: Effect of retraction on citations to prior work, including author position on retracted paper

Measure of Author Status	Absolute Status			Relative Status			Team Configuration		
	Total # of prior work	Total # of prior citations	h-index	Top 2 in Total # of prior work	Top2 in Total # of prior citations	Top2 in h-index	Total # of prior work	Total # of prior citations	h-index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated*Post(t>=1)	-0.136**	-0.135**	-0.156**	-0.213***	-0.191**	-0.196**	-0.055	-0.095	-0.017
	(0.068)	(0.063)	(0.065)	(0.079)	(0.082)	(0.081)	(0.081)	(0.104)	(0.075)
Author Status*Treated*Post(t>=1)	0.323	0.081**	0.208*	0.128***	0.103*	0.108*			
	(0.322)	(0.035)	(0.111)	(0.046)	(0.057)	(0.055)			
Self is eminent and Co-author is ordinary							-0.024	0.001	-0.049
*Treated*Post(t>=1)							(0.062)	(0.091)	(0.063)
Self is ordinary and Co-author is eminent							-0.124*	-0.124	-0.159*
*Treated*Post(t>=1)							(0.070)	(0.096)	(0.083)
Self is ordinary and Co-author is ordinary							-0.055	0.016	-0.091
*Treated*Post(t>=1)							(0.064)	(0.088)	(0.057)
Middle Author*Treated*Post(t>=1)	0.015	0.003	-0.001	0.002	0.002	0.001	0.007	0.006	0.001
	(0.080)	(0.076)	(0.078)	(0.077)	(0.074)	(0.077)	(0.077)	(0.078)	(0.078)
Last Author*Treated*Post(t>=1)	0.051	0.042	0.037	0.052	0.052	0.053	0.053	0.050	0.032
	(0.074)	(0.070)	(0.070)	(0.074)	(0.070)	(0.073)	(0.072)	(0.074)	(0.071)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	419,239	419,239	419,239	419,239	419,239	419,239	419,239	419,239	419,239
Number of papers	34,562	34,562	34,562	34,562	34,562	34,562	34,562	34,562	34,562

For interpreting regression coefficients in Columns (1)-(3), the total number of prior papers is measured in 1,000s; the total number of prior citations in measured in 10,000s and the h-index is measured in 100s. When author reputation is measured in relative terms, please look at the notes in Table 3 and 4. All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. All regressions include all one-way and two-way interactions terms; we do not report those coefficients for brevity. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table 8: Effect of retraction on citations to prior work, including author reputation at the time of publishing the retracted paper

Author Status measures	Full Sample			Ordinary Authors at Publishing		
	=1 if total # of prior work is in top 10% (1)	=1 if total # of prior citations is in top 10% (2)	=1 if h-index is in top 10% (3)	=1 if total # of prior work is in top 10% (4)	=1 if total # of prior citations is in top 10% (5)	=1 if h-index is in top 10% (6)
Treated*Post(t>=1)	-0.098** (0.041)	-0.086** (0.040)	-0.105** (0.042)	-0.097** (0.041)	-0.082** (0.040)	-0.105** (0.043)
Author Status at time of retraction *Treated*Post(t>=1)	0.180** (0.080)	-0.03 (0.084)	0.091* (0.047)	0.194** (0.082)	-0.054 (0.104)	0.106** (0.052)
Author Status at time of publication *Treated*Post(t>=1)	-0.125 (0.079)	0.065 (0.065)	-0.018 (0.043)			
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y
Observations	419,239	419,239	419,239	182,967	204,801	198,182
Number of papers	34,562	34,562	34,562	17,702	19,251	18,922

An author is defined ordinary at publishing if her absolute status measure fell out of the top 10 percentile of all treated authors at the time of publishing the retracted paper. Author status refers to the absolute status of a treated author if this author falls into top 10 percentile of all treated authors at the time of retraction. All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. All regressions include all one-way and two-way interactions terms; we do not report those coefficients for brevity. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Figure A1: Matching quality of control papers

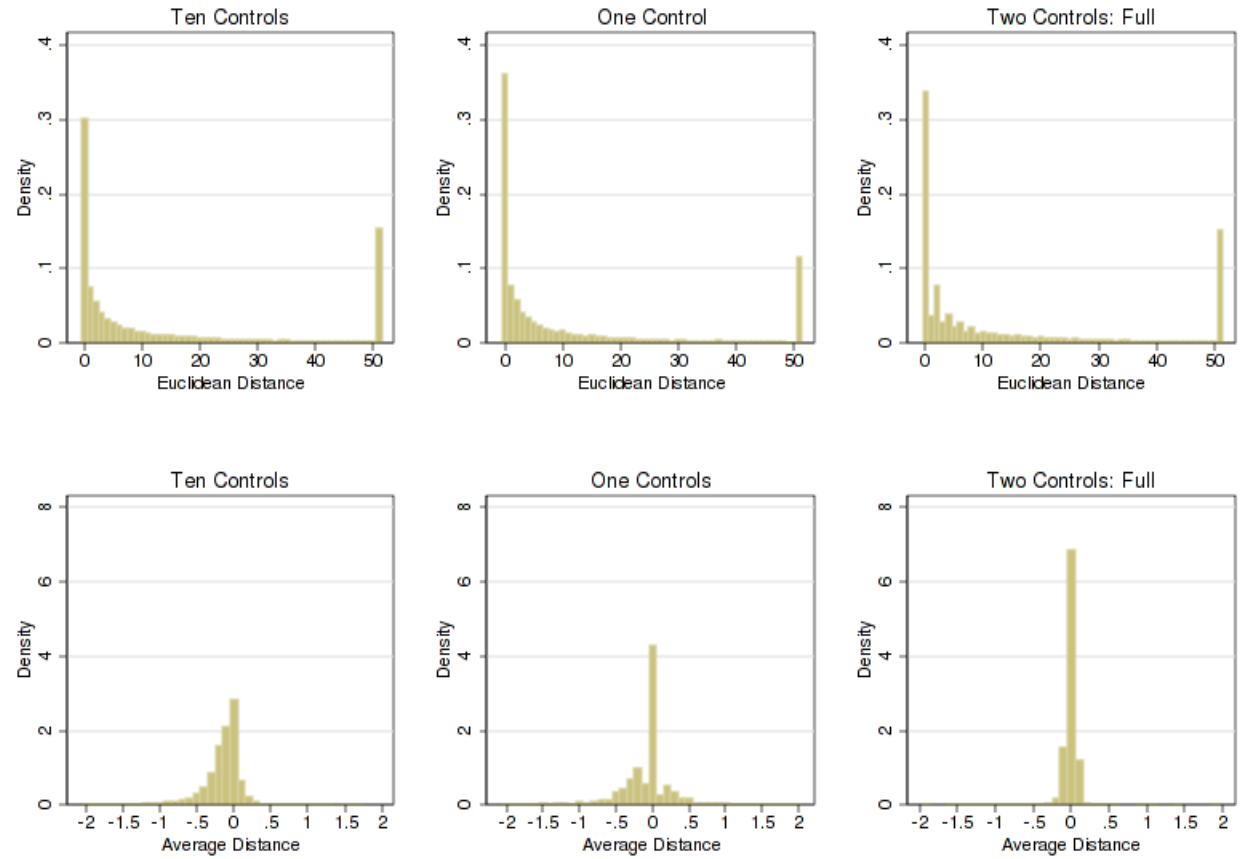



Figure A2: Distribution of h-index per treated author at the time of retraction



Note: we pool authors with an h-index greater than 80 at 80 in this figure.

Figure A3: citation life cycle of control papers




Table A1: Effect of retraction on log of citations to prior work, OLS

Measure of Author Status	Absolute Status			Relative Status			Team Configuration		
	Total # of prior work (1)	Total # of prior citations (2)	h-index (3)	Top 2 in Total # of prior work (4)	Top2 in Total # of prior (5)	Top2 in h-index (6)	Total # of prior work (7)	Total # of prior citations (8)	h-index (9)
Treated*Post(t>=1)	-0.061** (0.028)	-0.071*** (0.025)	-0.083** (0.033)	-0.129*** (0.045)	-0.118** (0.046)	-0.116** (0.046)	-0.034 (0.039)	-0.020 (0.034)	0.004 (0.034)
Author Status*Treated*Post(t>=1)	0.195 (0.204)	0.064** (0.032)	0.138 (0.086)	0.098** (0.042)	0.086* (0.044)	0.084* (0.043)			
Self is eminent and Co-author is ordinary *Treated*Post(t>=1)							0.007 (0.049)	-0.022 (0.048)	-0.040 (0.048)
Self is ordinary and Co-author is eminent *Treated*Post(t>=1)							-0.087 (0.061)	-0.124* (0.066)	-0.129* (0.074)
Self is ordinary and Co-author is ordinary *Treated*Post(t>=1)							0.005 (0.047)	(0.005) (0.044)	-0.045 (0.043)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	242,640	242,640	242,640	242,640	242,640	242,640	242,640	242,640	242,640
R-squared	0.268	0.268	0.268	0.268	0.268	0.268	0.268	0.268	0.268
Number of papers	34,562	34,562	34,562	34,562	34,562	34,562	34,562	34,562	34,562

All regressions are now ordinary least squares, with errors clustered by each retraction event. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table A2: Effect of retraction on citation to prior work, clustering by treated paper-control group

Measure of Author Status	Absolute Status			Relative Status			Team Configuration		
	Total # of prior work	Total # of prior citations	h-index	Top 2 in Total # of prior work	Top2 in Total # of prior citations	Top2 in h-index	Total # of prior work	Total # of prior citations	h-index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated*Post(t>=1)	-0.102***	-0.109***	-0.135***	-0.175***	-0.151***	-0.154***	-0.016	-0.059	0.009
	(0.029)	(0.025)	(0.034)	(0.041)	(0.041)	(0.041)	(0.034)	(0.041)	(0.032)
Author Status*Treated*Post(t>=1)	0.349*	0.083***	0.214***	0.121***	0.095**	0.097**			
	(0.182)	(0.026)	(0.077)	(0.041)	(0.040)	(0.040)			
Self is eminent and Co-author is ordinary							-0.029	-0.002	-0.056
*Treated*Post(t>=1)							(0.042)	(0.049)	(0.042)
Self is ordinary and Co-author is eminent							-0.123***	-0.126**	-0.165***
*Treated*Post(t>=1)							(0.045)	(0.055)	(0.051)
Self is ordinary and Co-author is ordinary							-0.063	0.009	-0.101**
*Treated*Post(t>=1)							(0.050)	(0.054)	(0.046)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	419,239	419,239	419,239	419,239	419,239	419,239	419,239	419,239	419,239
Number of papers	34,562	34,562	34,562	34,562	34,562	34,562	34,562	34,562	34,562

All regressions report coefficients from maximum likelihood estimation of a Poisson count model, but with errors now clustered by each treated paper control group. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table A3: Effect of retraction on citation of prior work, excluding treated papers that had zero citation in the year before retraction

Measure of Author Status	Absolute Status			Relative Status			Team Configuration		
	Total # of prior work	Total # of prior citations	h-index	Top 2 in Total # of prior work	Top2 in Total # of prior citations	Top2 in h-index	Total # of prior work	Total # of prior citations	h-index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated*Post(t>=1)	-0.106**	-0.111***	-0.140***	-0.175***	-0.159***	-0.155***	-0.009	-0.061	0.011
	(0.044)	(0.036)	(0.048)	(0.044)	(0.056)	(0.055)	(0.031)	(0.078)	(0.028)
Author Status*Treated*Post(t>=1)	0.390	0.084**	0.230**	0.121***	0.104*	0.099*			
	(0.324)	(0.035)	(0.109)	(0.045)	(0.058)	(0.056)			
Self is eminent and Co-author is ordinary *Treated*Post(t>=1)							-0.035	0.006	-0.053
							(0.057)	(0.093)	(0.058)
Self is ordinary and Co-author is eminent *Treated*Post(t>=1)							-0.142**	-0.129	-0.174**
							(0.062)	(0.096)	(0.079)
Self is ordinary and Co-author is ordinary *Treated*Post(t>=1)							-0.070	0.010	-0.105*
							(0.062)	(0.090)	(0.058)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	247,686	247,686	247,686	247,686	247,686	247,686	247,686	247,686	247,686
Number of papers	23,814	23,814	23,814	23,814	23,814	23,814	23,814	23,814	23,814

All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table A4: Effect of retraction on citation of prior work, excluding treated papers published within three years before retraction

Measure of Author Status	Absolute Status			Relative Status			Team Configuration		
	Total # of prior work (1)	Total # of prior citations (2)	h-index (3)	Top 2 in Total # of prior work (4)	Top2 in Total # of prior (5)	Top2 in h-index (6)	Total # of prior work (7)	Total # of prior citations (8)	h-index (9)
Treated*Post(t>=1)	-0.136*** (0.050)	-0.145*** (0.042)	-0.177*** (0.058)	-0.247*** (0.057)	-0.218*** (0.068)	-0.206*** (0.068)	-0.060 (0.055)	-0.077 (0.073)	-0.025 (0.044)
Author Status*Treated*Post(t>=1)	0.451 (0.365)	0.101** (0.042)	0.264** (0.131)	0.174*** (0.059)	0.142** (0.071)	0.128* (0.070)			
Self is eminent and Co-author is ordinary *Treated*Post(t>=1)							0.005 (0.076)	-0.002 (0.093)	-0.038 (0.071)
Self is ordinary and Co-author is eminent *Treated*Post(t>=1)							-0.143 (0.088)	-0.182* (0.105)	-0.210** (0.098)
Self is ordinary and Co-author is ordinary *Treated*Post(t>=1)							-0.039 (0.082)	0.010 (0.094)	-0.082 (0.072)
Author-Paper Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year Since Publication Dummies	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	359,273	359,273	359,273	359,273	359,273	359,273	359,273	359,273	359,273
Number of papers	25,187	25,187	25,187	25,187	25,187	25,187	25,187	25,187	25,187

All regressions report coefficients from maximum likelihood estimation of a Poisson count model, errors clustered by each retraction event. Standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Table A5: A Placebo Test: Do papers share the same citation patterns in the first several years have similar citation patterns later?

	Team Average (authors with prior)	Team Average (all authors)
Post($t \geq 1$)	0.873*** (0.188)	0.867*** (0.185)
Team Status*Post($t \geq 1$)	-0.014 (0.013)	-0.017 (0.017)

We conduct a placebo test by randomly sampling 500 pairs of clean (i.e., non-retracted) papers from our control sample. By construction, each pair has similar citation patterns prior to the (pseudo) retraction date. We next determine the author status for each control paper and further calculate the average author status among each paper's authors. We then examine whether higher status teams have different citation paths after the (pseudo) retraction event year for that pair. As can be seen from the interaction term in the table, the status measure has no predictive power for future citations. In other words, when two clean papers share similar citation patterns in the early stage, author status does not affect their citations in the later stage. Hence our control matches appear adequate to capture counterfactual citation paths, regardless of team status.