Essays on Value Creation, Disintermediation, and Platform-Based Strategies

Dissertation Advisors: Yuan Gu
Professors Feng Zhu, Shane Greenstein, Marco Iansiti

Essays on Value Creation, Disintermediation, and Platform-Based Strategies

Abstract

As firms increasingly rely on data and algorithms to create value, new strategic management challenges arise which lead to revenue loss and user attrition. My dissertation examines the strategic and managerial puzzles that hinder digital firms' value creation, from the perspectives of online community design, disintermediation, and platform-based strategies. The first chapter focuses on platform design and user segregation. The second and third chapters focus on disintermediation and its relationship to user trust and technological development. This dissertation aims to direct attention to strategic dilemmas that firms face as they embrace digital business models and provide theoretical and managerial contributions to platform studies in strategic management literature.

Keywords: Online platform, disintermediation, value creation, strategic management

Page intentionally left blank

Table of Contents

INTRODUCTION

Increasingly, organizations are encompassing a digital core – they are relying on data, information technology, and algorithms to create and capture value. In 2010, algorithm-based digital marketplaces contributed an estimated 34% of the US gross domestic product (Spulber 2011). Moreover, with the ongoing hit by COVID-19 on the global economy in 2020, firms are faced with the urgent need to adapt to an online model of businesses and operations more than ever. Such a switch to digital business models requires a new approach to strategy, while at the same time leads to unforeseen puzzles and challenges for managers.

My dissertation examines the strategic and management puzzles in any business model that encompasses a digital core, with a specific focus on two-sided platforms. It investigates value creation from the perspectives of online community design, disintermediation, and platform-based strategies. Do online communities lead to social segregation and a greater divide? Why do intermediaries capture less value as they provide better matches? How should platform owners better design the functional and technological features to facilitate user participation and avoid revenue loss? The three chapters in this dissertation, described below, seek to provide empirical and managerial insights to these questions.

The first chapter focuses on platform design and user segregation. Specifically, it directs attention toward platform features that facilitate shifts in ideology and composition among the online crowd, which may enable mutual dialogues or potentially exacerbate social segregation that hinders the platform's value creation.

The study in Chapter I, titled "Ideology and Composition among an Online Crowd: Evidence from Wikipedians," identifies the trend that participants on Wikipedia are converging and becoming more politically moderate over time. Two factors cause this shift toward moderation in collective opinion: either biased contributors contribute less, shifting the composition of participants, or biased contributors moderate their own views. By analyzing 2,887,140 participants and 10,878,391 contributions, the paper shows that participants tend to contribute to articles with slants that are the opposite of their existing political views. These contributors become more neutral when they encounter more extreme content of the opposite slant or when they receive more pushback from others. Extreme contributors may also leave the community after interacting with opposite-slant content and after receiving pushback. In sum, it sends an optimistic message: online communities can facilitate mutual understanding, and participants tend to become more politically neutral.

The second and third chapters focus on disintermediation, a strategic management problem that leads to major revenue loss for two-sided marketplaces. Increasingly, platforms are facing the risk of disintermediation, where customers and service providers circumvent the platform to transact directly. This hinders the platform's captured value and damages its business model. For platform owners, I highlight factors that can potentially lead to disintermediation, including direct trust between users and the development of communication technology.

In Chapter II, "Trust and Disintermediation: Evidence from an Online Freelancing Marketplace," I describe a puzzle regarding a platform's value creation and value capture: as a platform creates value by building trust between two sides of a market, it becomes increasingly difficult to capture that value due to disintermediation. Using a large-scale randomized controlled experiment, the paper finds that, while improved user trust increases the likelihood of high-quality service providers being hired, it also increases disintermediation, as users have less need for the platform in future transactions and, instead, transact directly to avoid paying the platform's fees. This offsets the revenue gains from hiring more high-quality sellers. These effects are stronger when both sides are geographically proximate to each other, when jobs are highly divisible, and when users have high ratings on the platform. Overall, this paper explores the theoretical contributions of disintermediation and provides managerial lessons on how platforms can mitigate the tension between trust-building and disintermediation.

Chapter III, titled "The Great Firewall of China and Marketplace Disintermediation," shows that the development of communication technology can also affect disintermediation. Based on a natural experiment that blocks third-party communication technologies outside the platform, I find that restricting alternative communication technology incentivizes users to communicate within the platform, thus significantly reducing the focal platform's disintermediation by 20.4% and leading to increases in the average transaction hours and fees. This effect is greater for transactions that are communication-intensive or posted by personal customers rather than business customers. This insight suggests that platforms should continue to improve their communication tools in order to create and retain a technological bottleneck. This strategy not only adds value to communication but also mitigates disintermediation.

My dissertation is among the first to investigate value creation from the perspectives of online community design, disintermediation, and platform-based strategies, using methods including randomized controlled trials, natural experiments, and statistical and textual analysis techniques on large-scale datasets. While the first chapter highlights platform features that facilitate shifts in ideology and composition among the online crowd, the key message is that platform design can be used to facilitate dialogues in the digital sphere, and that managers who aspire to produce content "from all sides" should let the most biased contributors leave the collective conversation. The second and third chapters add new empirical and managerial evidence to the research stream on disintermediation and value creation, as few studies in this literature explore the determinants of disintermediation as the mechanism of limiting value capture. To conclude, I hope that my dissertation directs attention to novel puzzles in platform-based strategies and provide theoretical and managerial contributions to platform studies in the field of strategic management.

CHPATER I:

Ideology and Composition among an Online Crowd:

Evidence from Wikipedians

I.1. Introduction

The growth of online communities that blur the boundaries between readers and writers has upended our understanding on the generation and consumption of online content. Online communities bring together participants from disparate traditions, with different methods of expression, different cultural and historical opinion foundations, and, potentially, different facts (e.g., Arazy et al. 2011; Ransbotham and Kane 2011; Kane et al. 2014; Gallus 2016).

Despite the diversity of opinions, and sometimes due to it, the composition and opinions of participants evolve as they interact with alternative content and points of view other than their own. A crowd's opinion reflects the aggregation of participants' opinions. Hence, at any point in time and over time, a crowd renders its opinion accordingly. Although a number of studies have sought to understand the bias in a crowd's opinion or, broadly, the limit of collective intelligence (e.g., Galton 1907; Shankland 2003; Antweiler and Frank 2004; Lemos 2004; Surowiecki 2004; Giles 2005; Chesney 2006; Rajagopalan et al. 2011; Mollick and Nanda 2015; Greenstein and Zhu 2018), little research has examined the manner by which the behavior of participants influences or is influenced by the bias of a crowd. Understanding this question helps shed lights on whether managers of such communities should intervene. It also informs managers on how to design effective rules and algorithms to steer interactions in ways that reduce a crowd's bias.

This study seeks to answer this question by measuring participants' actions and viewpoints. Our evidence comes from one of the longest-running online conversations on Wikipedia. We trace all the participation on 66,389 English language articles about US political topics from the start of Wikipedia in 2001 to January 2011. These articles received more than 10 million edits from 2,887,140 unique contributors. We follow Greenstein and Zhu (2018)'s approach to employ an adaptation of the method developed by Gentzkow and Shapiro (2010) for rating newspaper editorials. In these ratings, *slant* denotes the degree of opinion along a continuous yardstick. It can take on extreme degrees of red (e.g., Republican), extreme degrees of blue (e.g., Democrat), and all the shades of purple in between. *Bias* is the absolute value from the zero point of this yardstick and thus denotes the strength of the opinion. We then use these measures to characterize the evolution of the bias and slant of each participant opinion over time. We also gain insights into which experience prompts biased participants to stay or leave and which experiences induce them to maintain or change their opinion and, consequently, how such changes affect a crowd's bias.

Contributor behavior on Wikipedia tends to move toward less biased and less segregated conversations on most topics, consistent with Wikipedia's aspiration to present a neutral point of view (NPOV) in its content, which is succinctly summarized as "Assert facts, including facts about opinions—but don't assert opinions themselves."[1] Although considerable heterogeneity is found in behaviors, more Wikipedia contributors participate in unsegregated than segregated conversations. For example, a slanted contributor is on average 8% more likely to edit an article with an opposite slant than an article with the same slant. This tendency is pervasive.

We find that biased contributors moderate their own views as they encounter extreme content of the opposite slant or receive pushback from other contributors. Moreover, the composition of the existing contributors changes. We do not find evidence of a major change in the composition of new participants but do find evidence that more biased contributors exit sooner. Exposure to extreme opposite views is associated with a higher likelihood of withdrawal from participation. Furthermore, exit is the most significant driver of Wikipedia's bias. Simulations suggest that exit is responsible for 80.25%–90.98% of the decline in the slant.

[1] Source: https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view, accessed November 2018.

We examine a special circumstance, *mass edits*, where an article attracts an unusually high number of contributions in one day due to a sudden social event or breaking news about the topic. Such events are plausibly caused by factors exogenous to the Wikipedia community. During mass edits, articles experience more flips in slant in one day—from extremely blue/red to extremely red/blue. Consequently, contributors during mass edits are 11.8% more likely to be exposed to extreme content of both slants. As a result, contributors involved in mass edits demonstrate significantly faster reductions in slant than those involved in normal edits.

Our approach also allows us to analyze how fast someone changes his/her mind. For example, our estimates suggest that extreme Republicans take one year longer to become regular providers of neutral content than extreme Democrats. We trace this distinction to differences in the topics in which Democrats and Republican contributors participate in.

The findings offer important implications on the management of online communities. Past research typically focuses on various levers, such as social, psychological, and economical, available to managers of online communities, which they can use to maximize participation or minimize churn (e.g., Lerner and Tirole 2002; Wasko and Faraj 2005; Bagozzi and Dholakia 2006; Jeppesen and Frederiksen 2006; Moon and Sproull 2008; Nambisan and Baron 2010; Zhang and Zhu 2011; Gallus 2016; Nagle 2018). We find that one of the key mechanisms on how Wikipedia achieves an NPOV is by letting contributors with extreme viewpoints leave the communities. As long as new contributors continue to arrive, we see less reason for community managers to maximize participation or be overly concerned about the exit of participants. For other community managers, if they spire to draw on multiple opinions and achieve a balance between them in their online communities, they must insist that contributors also aspire to that goal and actively discourage participation from those who maintain extreme points of view.

We also identify a key feature of Wikipedia that facilitates the convergence to neutrality, that is, contributors are frequently exposed to the content of opposite slants. In practice, however, various communities often design algorithms to expose their contributors to content that aligns with their preferences. Although this strategy maximizes participants' satisfaction, as shown in our research, such practices are harmful in building a less-polarized and unbiased crowd.

The importance of Wikipedia in the modern society makes understanding its production interesting in its own right. Most reference information has moved online, and these online sources have displaced other sources of information in every developed country. Wikipedia is the top 20 site in several developed countries and, by far, the most popular and referenced online repository of comprehensive information in the developed world. The English language version of Wikipedia has received over eight billion page views per month and over 500 million unique visitors per month.[2] Many firms also utilize Wikipedia as an input. Amazon (Alexa), YouTube, and Google (search), among others, use Wikipedia as a free source for neutral "facts" and as an unrestricted source for vocabulary in different languages.[3]

I.2. Relationship with Prior Work

Considerable research has examined the property of online crowds. Although some studies show that collective decision making can be more accurate than experts' decision making (e.g., Antweiler and Frank 2004; Lemos 2004; Surowiecki 2004; Giles 2005; Rajagopalan et al. 2011), others find that a crowd can be more biased (e.g., McPherson et al. 2001; Sunstein 2001; Rector

---

[2] "Wikipedia vs. the small screen." http://www.nytimes.com/2014/02/10/technology/wikipedia-vs-the-small-screen.html?_r=1, accessed June 2019.

[3] See, e.g., "YouTube May Add to the Burdens of Humble Wikipedia," https://www.nytimes.com/2018/03/19/business/media/youtube-wikipedia.html, accessed June 2019.

2008; Gentzkow and Shapiro 2011; Park et al. 2013; Bail et al. 2018; Greenstein and Zhu 2018). For instance, Gentzkow and Shapiro (2011) find biases in online conversations about political content and other topics higher than the segregation of offline news consumptions**.** Greenstein and Zhu (2018) show that Wikipedia articles are on average more biased than those in Britannica, an encyclopedia authored by experts. Several studies have proposed new approaches to aggregate opinions from the crowds to minimize bias (e.g., Fan et al. 2005; Muchnik et al. 2013; Prelec et al. 2017; Wang et al. 2017). Although these studies acknowledge that the bias of a crowd reflects the aggregate of individual participants, to the best of our knowledge, no studies have tracked long-run changes in how extremists (moderates) participate and whether they change their expression to more moderate (extreme) views. We think that this gap arises, in part, because it is rare to observe an online crowd over a long time period and a wide array of topics. This study has an example of such case in Wikipedia, analyzing almost a decade of participation.

This concern becomes more relevant when participants confront *contested knowledge*—defined as topics involving subjective, unverifiable, or controversial information. Many observers are worried about the emergence of *segregated* conversations in the presence of contested knowledge in online crowds. Segregated conversation may become an "echo chamber" (EC) of like-minded views (e.g., Sunstein 2001; Van Alstyne and Brynjolfsson 2005; Carr 2008; Quattrociocchi et al. 2016; Shore et al. 2016; Sun et al. 2017). The opposite behavior, an *unsegregated* conversation, involves contributors with diverse ideas and opposing views (Benkler 2006). Many unsegregated conversations bring varying perspectives into a common view by accelerating a confrontation or discourse between contradictory facts and ideas. Generally, segregated conversations are blamed for many undesirable outcomes.[4]

In the case of contested knowledge in Wikipedia, prior research (Greenstein and Zhu 2012, 2016, 2018) shows that the slant and bias of content evolve and bias in Wikipedia articles slowly declines over time. Prior research does not identify the underlying mechanism other than to affiliate it with more revision. The shift in the slant and bias can be caused by many factors, such as the arrival of moderate contributors, withdrawal of extremists, or changes in contributors' own viewpoints. Without examining actual participants' behavior, it is difficult to draw the right managerial implications for community managers. For instance, Greenstein and Zhu (2018) find that bias in Wikipedia articles tends to slowly decrease with more revisions. One may infer from this result that community managers need to encourage more participation to reduce bias faster. Our research shows that the situation is not just about the number of contributions but about the identity of the contributors. More contributions reduce bias because over time these contributions mostly come from moderate contributors. Fundamentally, the change in the composition and ideology of the crowd drives the bias reduction. Hence, different from prior research, our study suggests that community managers should devote efforts into designing processes to encourage extremists to leave or convert themselves into more moderate contributors rather than to maximize participation.

---

[4] As already implied above, concerns about the health and tenor of political conversations have motivated prior works (e.g., Sunstein 2001; Carr 2008; Lawrence et al. 2010; Gentzkow and Shapiro 2011; Greenstein and Zhu 2012, 2016; Shore et al. 2016; Boxell et al. 2017). Closer to our study, Gentzkow and Shapiro (2011) focus on online conversations about political content and other topics, and Gentzkow and Shapiro (2010) start from the premise that ideological tendencies appear in the language of speakers. Segregation can facilitate the radicalization of some individuals and groups (Purdy 2015). See, for example, http://www.vice.com/read/we-asked-an-expert-how-social-media-can-help-radicalize-terrorists and http://www.rand.org/randeurope/research/projects/internet-and-radicalisation.html, accessed June 2017. Segregated conversation can also discourage interracial friendships, disconnect different social segments, and stimulate social isolation. In traditional media, ideological biases in news content affect the political language (e.g., DellaVigna and Kaplan 2007; Stone 2009; Chiang and Knight 2011; Durante and Knight 2012).

In addition, prior studies show that people's beliefs become more reinforced when they encounter information that is aligned with their prior beliefs (e.g., Van Alstyne and Brynjolfsson 2005; Gilbert et al. 2009; Bakshy et al. 2015; Lee et al. 2015; Garimella et al. 2018). However, the manner by which people react to opinions that differ from their beliefs is unclear. For example, studies find that people may demonstrate a "pushback": they refute evidence that has a contrary effect on belief (e.g., Nyhan and Reifler 2010; Wood and Porter 2019). In some rare cases, we may observe a "backfire effect": given evidence against their beliefs, people can reject the evidence and believe even more strongly.[5] Our empirical results that many extreme contributors choose to leave after encountering opposite opinions provide support for the claim that changing people's beliefs is difficult. At the same time, a number of contributors become more moderate after encountering opposite opinions. Although such changes are responsible for a small fraction of overall bias reduction, the evidence restores our hope that communities, such as Wikipedia, can help reduce polarization in our society as they gradually work toward an NPOV.

This study is also related to the literature on platform design for user-generated content.[6] Although much of this literature has examined how algorithms (un)intentionally nudge user behavior in one direction or another and how they may produce unanticipated aggregate outcomes because they often seek to match content with a user's taste to maximize a user's satisfaction or to retain users, algorithms play no role in our setting. Wikipedia employs an architecture that gives

---

[5] See, for example, "The Backfire Effect," https://archives.cjr.org/behind_the_news/the_backfire_effect.php, accessed August 2019.

[6] Prior research has examined the importance of contributor motivation for a variety of tasks, such as software design, entrepreneurial finance, and engineering (e.g., Kogut and Metiu 2000, 2001; Rothaermel and Sugiyama 2001; Von Krogh and Von Hippel 2006; Chesbrough 2006; Roberts et al. 2006; Yang et al. 2009; Ramasubbu and Balan 2012; Ransbotham et al. 2012; Kane et al. 2014; Xu et al. 2015; Gallus 2016; Nagaraj 2017; Qiu and Kumar 2017). Most empirical studies have examined how online organizations aggregate contributions to solve collective problems (e.g., Kogut and Zander 1992; Lee and Cole 2003; Hargadon and Bechky 2006; Kuk 2006; Tucci and Villarroel 2007; Xu and Zhang 2009, 2013; Faraj et al. 2011; Ransbotham and Kane 2011; Afuah and Tucci 2012; Chen et al. 2012; Pierce 2012; Bassamboo et al. 2015).

participants considerable discretion in achieving platform-wide ideals and aspirations.[7] Other studies have examined segregation of participants in social networks, such as Twitter. Shore et al. (2016) study sharing links on Twitter and examine whether participants share with others who are like-minded. Bail et al. (2018) study a field experiment about following opinion leaders on Twitter and examine whether exposure to opposite viewpoints changes users' ideologies over time. This scenario is similar in our study, although the exposure in their study is monetary incentivized. Both Shore et al. (2016) and Bail et al. (2018) examine whether social interactions reinforce segregated conversation, but they reach different conclusions. We regard our setting as an opportunity to understand user behavior in the absence of algorithms and social networking features. Our study suggests that the platform at risk of losing users may actually provide the optimal solution.

I.3. Empirical Setting

Founded in 2001, Wikipedia positions itself as "the free encyclopedia that anyone can edit" or, in other words, as an online encyclopedia entirely written and edited through user contributions. Topics are divided into unique pages, and users can select any page to revise. It has become the world's largest "collective intelligence" experiment and one of the largest human projects ever to bring information into one source.

Contributions come from tens of millions of dedicated contributors who participate in an extensive set of formal and informal roles.[8] Some roles entail specific responsibilities in editing

[7] Similar with other online communities, Wikipedia has adopted explicit aspirations, rules, norms, policies (Forte et al. 2009; Jemielniak 2014; Schroeder et al. 2012), and quality assurance procedures (Stvilia et al. 2008), which shape contributors' behavior. Many online communities have adopted privilege access schemes that formally define roles (Arazy et al. 2015; Burke et al. 2008; Collier et al. 2008; Forte et al. 2012), and Wikipedia has performed this as well. This initiative has led to a myriad of coordination mechanisms (Kittur et al. 2007a; Kittur and Kraut 2008; Kittur et al. 2007b; Schroeder and Wagner 2012), social interactions (e.g., Halfaker et al. 2011; Forte et al. 2012), and behaviors aimed at conflict resolution (Arazy et al. 2011).

[8] See https://en.wikipedia.org/wiki/Wikipedia:User_access_levels, accessed June 2017.

tasks; however, the Wikimedia Foundation employs a limited set of people and does not generally command its volunteers. Instead, it develops mechanisms to govern the volunteer co-production process (Kane and Fichman 2009; Te'eni 2009; Zhang and Zhu 2011; Hill 2017). All these voluntary contributors are considered editors on Wikipedia. The organization relies on contributors to discover and fix passages that do not meet the site's content tenets. However, no central authority tells contributors how to allocate their editorial effort.

The reliance on volunteers has many benefits and drawbacks. Among the latter, there is a long-standing concern that interested parties attempt to rewrite Wikipedia to serve their own parochial interests. Despite the persistence of such concerns, little systematic evidence has pointed in one direction or another. The available evidence on conflicts suggests that contributors who frequently work together do not get into as many conflicts as those who do not, nor do their conflicts last as long (Piskorski and Gorbatâi 2017). Although such behavior can lead to edits from contributors with different points of view, no direct evidence shows that it leads to more content that finds compromises between opposite viewpoints.

Although the Wikipedia attempts to attract a large and diverse community of contributors, it also invites many slanted and biased views, and the openness of Wikipedia's production model (e.g., allowing anonymous contributions) is subject to sophisticated manipulations of content by interested parties. Hence, there is a widespread acceptance of the need for constant vigilance and review.

A key aspiration for all Wikipedia articles is an NPOV (e.g., Majchrzak 2009; Hill 2017). To achieve this goal, "conflicting opinions are presented next to one another, with all significant points of view represented" (Greenstein and Zhu 2012). When multiple contributors make inconsistent contributions, other contributors devote considerable time and effort debating whether the article's text portrays a topic from an NPOV. Because Wikipedia articles face few realistic limits regarding their number or size[9] (due to the absence of any significant storage costs or any binding material expense), conflicts can be addressed by adding more points of view to articles instead of eliminating them (e.g., Stvila et al. 2008). In general, most disputes are settled without interventions from Wikipedia administrators.[10]

I.4. Data and Summary Statistics

I.4.1. Measuring Contributor Slant and Bias

We extend the approach of Greenstein and Zhu (2018) to measure Wikipedia contributors' slants and biases. This approach relies on the modification of an existing method, developed by Gentzkow and Shapiro (2010), for measuring slants and biases in newspapers' political editorials.[11] For example, Gentzkow and Shapiro (2010) find that Democratic representatives are more likely to use phrases, such as "war in Iraq," "civil rights," and "trade deficit," whereas Republican representatives are more likely to use phrases, such as "economic growth," "illegal immigration," and "border security."[12] Similarly, we compute an index for the slant of each article

[9] Over time, a de facto norm has been developed that tends to keep most articles under 6,000 to 8,000 words. This guideline has arisen as editorial teams have debated and discussed the article length necessary to address the topic of the page. Of course, some articles grow to enormous length, and editor contributors tend to reduce this length by splitting them into sub-topics. A prior work (Greenstein and Zhu 2016) finds that the average Wikipedia article is shorter than this norm (just over 4,000 words), but the sample includes a few longer articles (the longest is over 20,000 words).

[10] Similar with all matters at Wikipedia, contributors have discretion to settle disputes on their own. The organization offers a set of norms for the dispute resolution processes, which can be quite elaborate, including the three-revert edit war rule and rules for the intervention of arbitration committees and mediation committees. Administrators can also decide to freeze a contentious article.

[11] Gentzkow and Shapiro (2010) characterize how newspapers also use such phrases to speak to constituents who favor one political approach over another.

[12] Several studies have applied their approach in analyzing political biases in the online and offline content (e.g., Greenstein and Zhu 2012; Jelveh et. al. 2014; Shore et al. 2016). In addition, although Budak et al. (2016) use alternative approaches to measure ideological positions of news outlets, their results are consistent with those of Gentzkow and Shapiro (2010).

from each source, tracking whether articles employ words or phrases that appear to slant toward either Democrats or Republicans.[13]

Initially, we assume that a contributor's slant is constant throughout the years and define a contributor's slant as the average slant of all the contributions that the person made in our sample. Then, we allow a contributor's slant to evolve over time. The measure of a contributor's slant is computed based on the contributions in each year instead of throughout the sample period. A contributor's bias is the absolute value of the slant in both cases.

To construct our sample, we focus on broad and inclusive definitions of U.S. political topics, including all Wikipedia articles that include the keywords "Republican" or "Democrat." We start by gathering a list of 111,216 relevant entries from the online edition of Wikipedia on January 16, 2011. Eliminating the irrelevant articles and those concerning events in countries other than the U.S.[14] reduces our sample to 70,305. Our sample includes topics that are highly debated, such as abortion, gun control, foreign policy, and taxation, and less disputed ones relating to minor historical and political events and biographies of regional politicians. We then collect the revision history data from Wikipedia on January 16, 2011, which yields 2,891,877 unique contributors.

Our key dependent variable is *Contributor Slant*. This measure is developed in two steps. First, every article on Wikipedia has a revision history that, for every edit, records pre-edit and post-edit versions. We compute the slant index for the pre- and post-edit article versions, take the difference between the two, and use this difference as the *slant change* for an edit. We obtain the slant change

[13] An article's slant changes only when code phrases are added and/or dropped.

[14] The words "Democrat" and "Republican" do not appear exclusively in entries about U.S. politics. If a country name shows up in the title or category names, we then check whether the phrase "United States" or "America" shows up in the title or category names. If yes, then we keep this article. Otherwise, we search the text for "United States" or "America." We retain articles in which these phrases show up more than three times. This process allows us to keep articles on issues, such as the "Iraq war," but excluded articles related to political parties in non-U.S. countries.

of every edit. For sequential edits from the same contributor that happened consecutively and without anyone else editing between them, we treat the sequence of edits as one single edit.[15]

To analyze participant behaviors, we exclude the first version of all articles in our sample (or if the article has only one version, then the whole article) as we do not have a prior article slant and cannot observe the EC or non-EC effect for such contributions. We also exclude contributors who made more than 950 edits in any one year (top 0.01%), since these contributors could be bots that regularly maintain Wikipedia or contributors who created many articles when Wikipedia was first founded. These procedures reduce the number of observations in the sample to 9,487,164, the number of articles to 65,046, and the number of unique contributors to 2,886,795. We use this sample as the main analysis sample.

Next, we focus on individual contributors. We identify and measure the types of changes that each contributor makes to Wikipedia articles. We assign each edit to each contributor and assign a slant value for each edit. Under the assumption that every contributor has one fixed type of slant, we compute the *Contributor Slant* as the average value of the slant index of this contributor. A zero value of *Contributor Slant* means that the user's edits either contain a balanced set of Republican/Democratic words (weighted by their cardinal values) or do not include any of the slanted phrases. A negative or positive value of *Contributor Slant* means that the contributor is Democrat or Republican leaning, respectively. Accordingly, the absolute value of a contributor's slant is equal to the contributor's bias. In our sample, 92.6% of the contributors have a zero slant, whereas the remaining 225,000 contributors make at least one slanted contribution. As it turns out,

[15] These consecutive edits tend to be highly correlated, or they can be several parts of a complete contribution, such as where the contributors saved their work several times. As a robustness check, we exclude deletions from a contributor's edits if the deletion does not bring an article's slant from left/right leaning to right/left leaning or from less to more extreme. Accordingly, deleting biased content to make an article more neutral will not be considered a biased edit. Accordingly, all of our results still hold.

the majority (57.5%) of contributions to Wikipedia come from contributors with a measurable slant or bias.

Table 1: Distribution of Different Types of Contributors over Years

Year	Democrat Contributors	Core Democrat Contributors	Republican Contributors	Core Republican Contributors	Neutral Contributors	Core Neutral Contributors	Total # of Contributors Contributed in the Year
2001	26.4%	18.1%	20.0%	12.5%	53.6%	9.9%	800
2002	9.9%	7.5%	9.6%	7.4%	80.4%	17.6%	4,364
2003	8.5%	6.5%	8.8%	6.9%	82.6%	18.3%	14,951
2004	7.8%	5.7%	7.7%	5.9%	84.5%	17.3%	66,867
2005	7.0%	4.7%	6.7%	4.6%	86.3%	15.6%	242,121
2006	5.7%	3.6%	5.7%	3.6%	88.6%	14.7%	584,438
2007	5.3%	3.2%	5.2%	3.3%	89.5%	13.8%	706,195
2008	5.2%	3.1%	5.3%	3.2%	89.5%	13.9%	640,871
2009	4.7%	3.1%	4.7%	3.2%	90.5%	14.1%	526,255
2010	4.2%	2.8%	4.2%	2.9%	91.6%	13.2%	461,663
2011	9.5%	8.5%	10.8%	9.9%	79.6%	19.4%	26,886

Notes: "Democrat/Republican/Neutral Contributors" shows the percentage of contributors with negative/zero/positive *Contributor Slant* among all contributors who contribute in that year to the articles in our sample. "Core Democrat/Republican/Neutral Contributors" shows the percentage of that year's "Democrat/Republican/Neutral Contributors" whose total number of edits is distributed in the top 10% of all contributors' total number of edits. Final year, 2011, is sampled in January, which accounts for the low numbers in that year.

Table 1 presents the distribution of contributor types over a 10-year period. When computing the number of Democratic, Republican, and neutral contributors to Wikipedia each year, we count each contributor only once, even if the contributor contributes many times in a year. We summarize the distribution of the contributors' total number of edits over the 10 years in Figure 1. Our sample reflects the well-known skewness of contributions to Wikipedia. More than 75% of the contributors in our sample contributed only once in the entire 10-year period, whereas 97.5% of the contributors contributed fewer than 10 times, averaging less than one contribution per year. Only 1% of the contributors contributed more than 30 times in our sample. We also show the number of edits, number of contributors, and average number of edits per contributor by the contributors' years of experience in Figures 2–4, respectively. Although contributors with four to five years of experience comprise the large part of our sample in terms of the number of contributors and the total number of edits, the average number of edits per contributor does not vary much with years of experience, except for the 0.18% of contributors who joined in January 2011.[16]

We define a contributor as *core* if his or her total number of edits is distributed in the top 10% of all contributors' total number of edits, which in this case is equal to a total of no less than three contributions. Core contributors make 74% of the contributions in the entire sample. In other words, 84.2% of the edits for each article are from core contributors and 15.8% edits are from peripheral contributors. Furthermore, although the number of neutral contributors who contribute each year is more than 10 times the number of contributors who have a slant, the proportion of core contributors in the neutral slant group (15.9% for the full sample) is much smaller compared with that in the other two groups (63.8% and 65.5% for the full sample, respectively). In summary, slanted contributors are more core than neutral contributors, and much of the slanted content comes from contributors making many edits.

[16] Excluding this group of contributors does not change the qualitative results.

Figure 1: Distribution of Contributors' Total Number of Edits

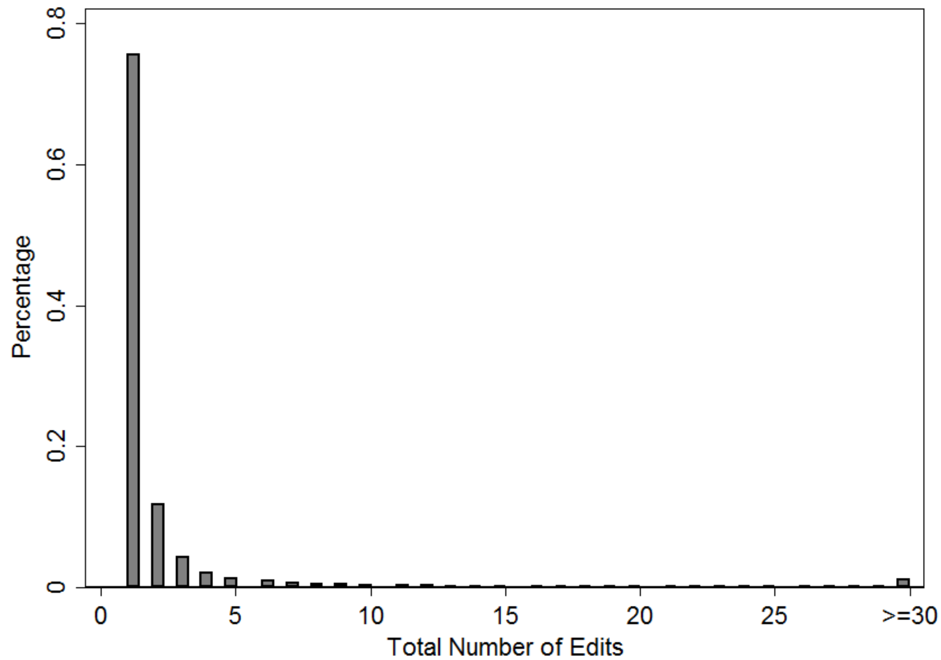

Figure 2: Distribution of All Edits in the Sample by Contributors' Years of Experience

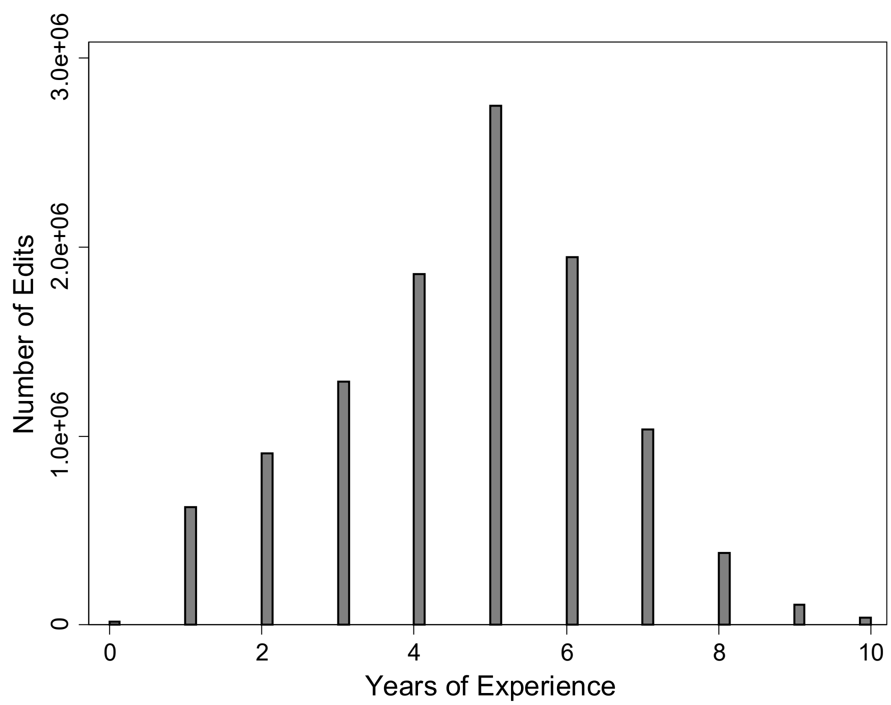

Figure 3: Distribution of Number of Contributors by Years of Experience

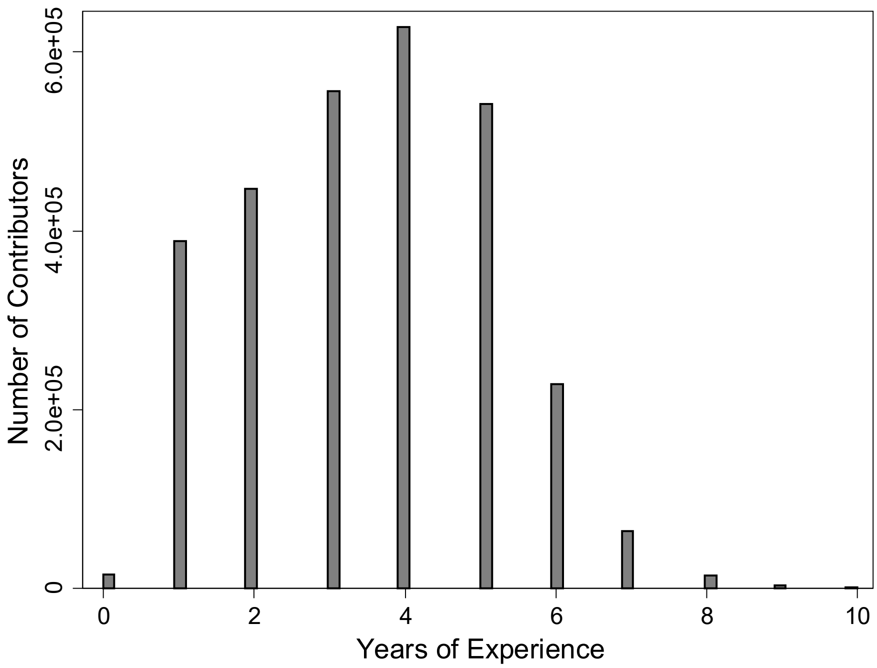

Figure 4: Distribution of Average Number of Edits per Contributor by Years of Experience

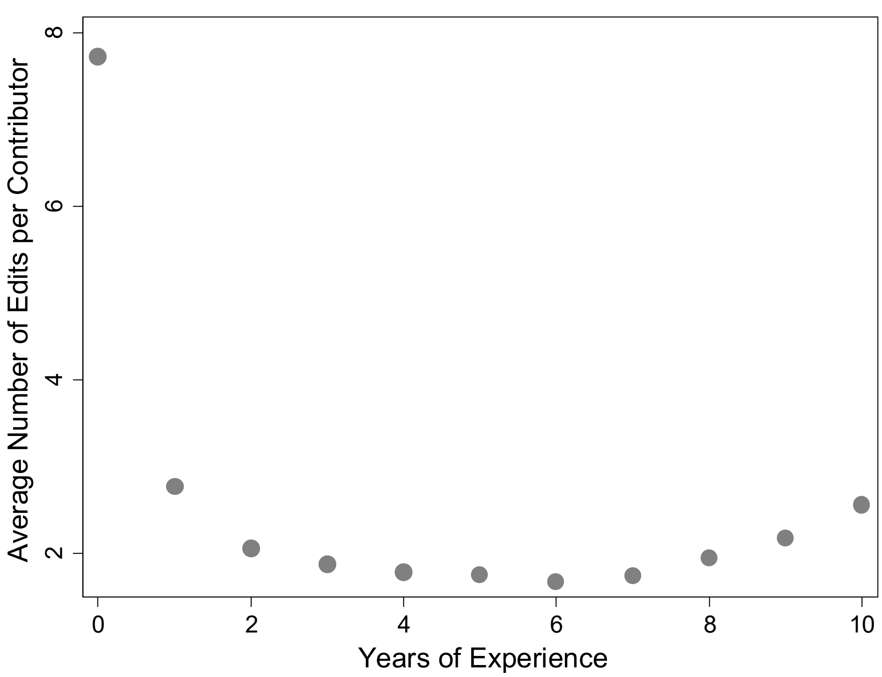

I.4.2. Composition or Ideological Shift?

A few simple graphs illustrate the evolution of contributors' biases. Figure 5 is a bar chart of the average bias of any contributor who contributed at least once in that given year. In addition, in the graphs, we do not plot the observation in 2001 because contributors in the founding year of Wikipedia tend to be different from contributors in later years, and they only account for 0.03% of the full sample.[17] The average bias of the contributors declines over the years.

Figure 5: Average Contributor Bias over the Years

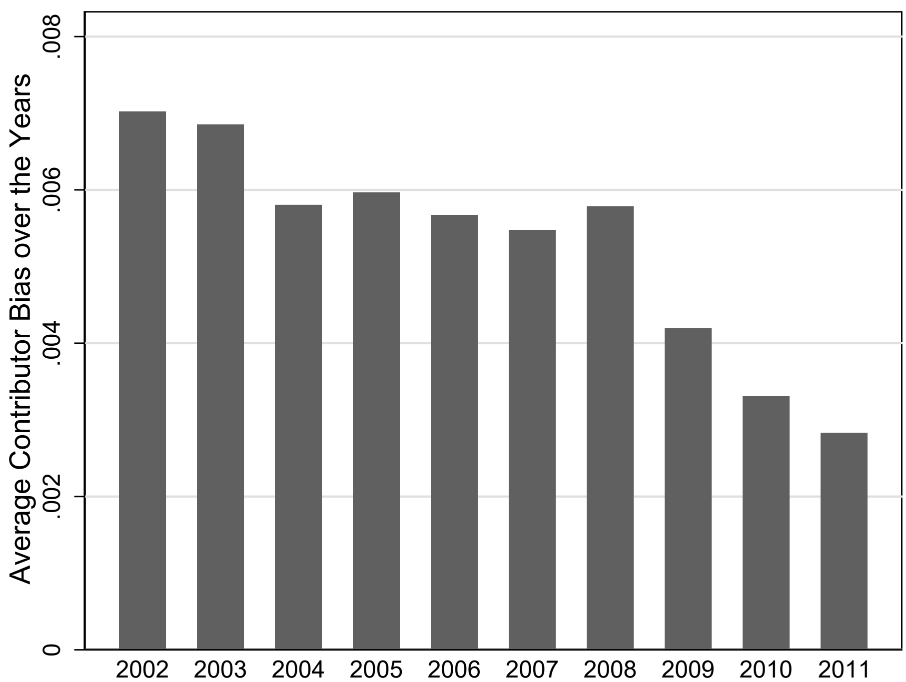

Two types of changes may have contributed to the decline in bias: a change in the composition of the contributors and/or an ideological shift. That is, new contributors with a moderate bias may join Wikipedia each year, and relatedly, the existing extreme contributors edit less over time or gradually stop participating. Alternatively, contributors can become less biased in their

[17] Excluding all contributors who joined in 2001 does not change the qualitative results.

contributions, in which case our assumption of a fixed slant for each contributor requires modification.

First, we consider the changes in the biases of people joining Wikipedia in different years. We compute the average bias of contributors entering in different years and plot the results in a bar chart in Figure 6. We do not observe an obvious pattern across years. In general, except in 2002, contributors who entered earlier are not systematically more biased compared with those who entered later.[18]

Figure 6: Vintage Effect for Contributors Entering in Different Years

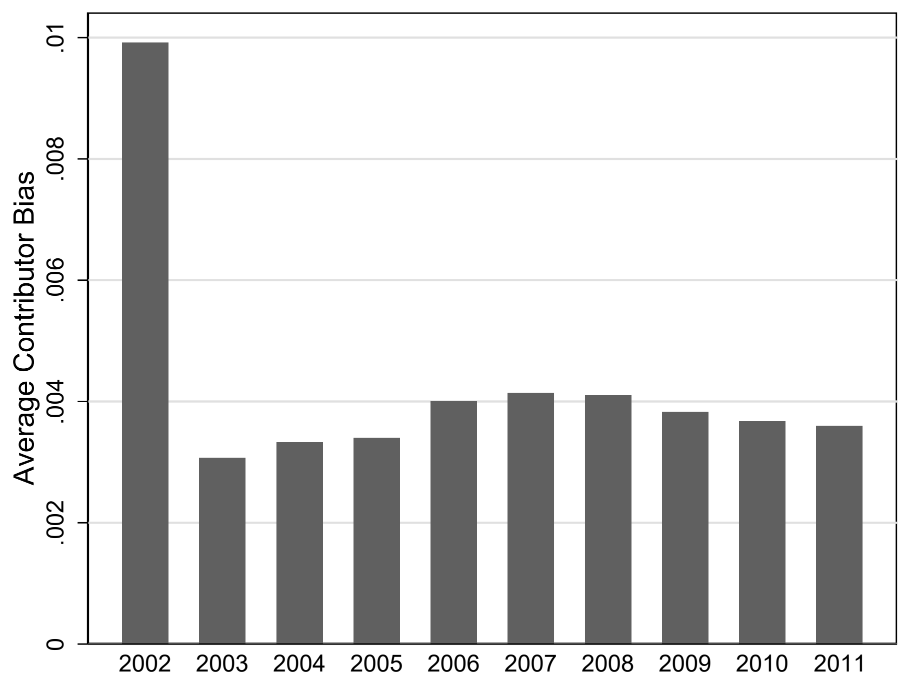

Next, we consider existing contributors on Wikipedia. Figure 7 displays the average number of contributions of the extreme contributors each year. A contributor is considered "extreme" if

[18] The contributors who entered in 2002, the second year of Wikipedia's existence, are only 0.13% of the full sample. Prior work has shown that the earliest Wikipedian's tended to be extreme Democrat-leaning contributors and compositition quickly changed to more moderate participants on average and years before Wikipedia became popular. See Greenstein and Zhu (2016).

his/her slant across all years is more than two standard deviations away from the center. Although the number of edits from the contributors' first year to their second year on Wikipedia seems to increase, this is followed by a declining pattern as they stay past two years. In other words, as these extreme contributors stay longer, they become less active over time. This pattern indicates the exit of extreme contributors as a potential explanation for the overall bias decline.

Figure 7: Average Number of Edits over Contributors' Years on Wikipedia, Extreme Contributors Only

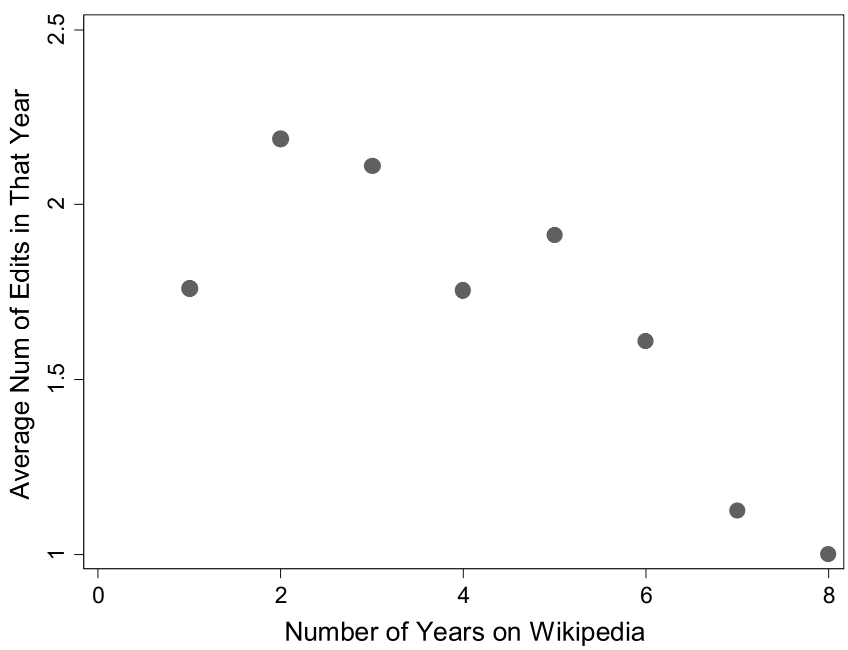

For extreme contributors who keep contributing over the years, we then ask whether we can observe a decline in the biases of their contributions. Figure 8 plots the average contributor bias each year for those who are considered "extreme." If we redefine the slant and bias each year (based on their changes in that year), then we observe a constant declining pattern in the biases from these extreme contributors. The pattern suggests that extreme editors become less slanted over time.

Figure 8: Average Contributor Bias Each Year over Contributors' Years on Wikipedia, Extreme Contributors Only

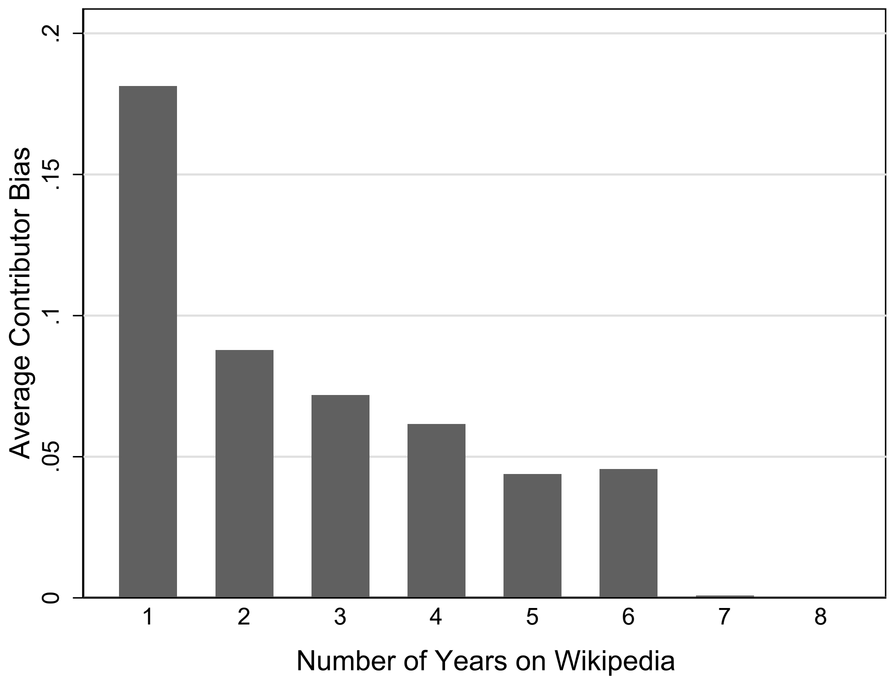

Notes: In the bar chart, the average bias in year seven is 0.0007 and in year eight is 0.

Overall, these graphs suggest a change in the composition of extreme contributors and their participation mostly due to the exit of participants and an ideological shift favoring more neutral contributions. The arrival of new participants with neutral views appears to play little role. Do these patterns survive more careful statistical scrutiny? What factors drive the observed behavior patterns? We next define variables in preparation for investigating these questions.

I.4.3. Variables

Contributor Slant assumes that contributors have the same slant over their lifetimes in our sample. We next define *Contributor Yearly Slant*, which divide contributors' edits by year, and for each year use the same calculation as for *Contributor Slant* (i.e., we compute the average slant change of all the edits that a contributor has made within a given year). If a contributor's numeric value for slant remains unchanged throughout the years, then his or her *Contributor Yearly Slant* is equal to *Contributor Slant*.[19] Relatedly, *Contributor Yearly Bias* is the average absolute value of the slants (i.e., the biases of a contributor's edits in a given year). The purpose of this variable is to better capture a contributor's bias.

Prior Article Slant denotes an article's slant before a given edit. This variable is essential for analyzing an article's (mis)match with a contributor's slant.

To measure a contributor's experience in interacting with different types of content, we count the contributor's number of edits in a given time period targeting extreme opposite-slant articles, divided by the contributor's total number of edits during that time, and label it as *Opposite-Slant Article Edits Fraction.* Similarly, the proportion of a contributor's edits targeting extreme same-slant articles is labeled as *Same-Slant Article Edits Fraction*, which captures the amount of extreme content that he or she has interacted with. For a contributor who only made neutral edits in a given year, that is, the contributor is considered neutral, *Opposite Slant Article Edits Fractions* and *Same Slant Article Edits Fractions* is equal to zero as either extreme-left or extreme-right content should be seen as "opposite to" or "same as" a neutral person's ideology.

Pushbacks may shape contributors. An extreme example of a pushback is the *revision war*, where a contributor's edit is immediately reverted by another contributor. In this case, the original contributor edits back the same contribution. We count the edits that a contributor makes during such revision wars, divided by the contributor's total number of edits, as *Revision War Edits Fraction*.

[19] If a contributor does not make any contribution in a given year, then his or her *Contributor Yearly Slant* has a missing value in that year.

Throughout the study, unobservable features of articles are a central concern. We add additional measures that may have attracted editors and otherwise had a spurious correlation with the slant or bias of an article. We measure the length of the articles using the number of words in an article prior to a certain edit and label it as *Prior Article Length.* We measure the number of the article's external references and label it as *Prior Refs*. Articles that are long may incorporate more viewpoints, which in turn tends to attract more contributors. In addition, Wikipedia requires citations from major third-party sources as references for its article content (often listed at the bottom of the page), so articles with more references are also more likely to incorporate more outside arguments or controversial views at the time.

In a similar vein, additional controls measure editors' unobservable features. One such variable is *Number of Edits*, which is the total number of *yearly* edits that the contributor has made on Wikipedia in a given year. Another variable is *Starting Number of Edits*, which is equal to the total number of edits that the contributor made in his or her first two years after joining Wikipedia.

To examine the determinants of composition, we define a dummy variable *Stay* at the contributor level. *Stay* is equal to 1 if the contributor made at least one edit in the last year in our sample; otherwise, it is equal to 0.[20]

Table 2 provides the summary statistics of all the variables used in our analysis. The average *Contributor Slant* in our sample is negatively close to zero, indicating that Democrat-leaning contributors are, on average, more slanted than Republican-leaning contributors. Moreover, the article versions in our sample exhibit similar absolute values of extreme slant on both ends. For control variables *Prior Article Length*, *Prior Refs*, and *Number of Edits*, a substantial variation is

[20] We did not construct *Stay* as a time-varying dependent variable by year because only 1.95% of the contributors in our sample was inactive (i.e., "exited") in one year and came back to edit again in another year.

found across article versions for each of the measures, and we use the logarithm of these control variables in our models because they are highly skewed.[21]

Table 2: Summary Statistics of Variables Used in the Main Analyses

Variable	Unit	Mean	Std. dev.	Median	Min	Max
For Participation Pattern Analysis:						
Contributor Slant	Contribution	-0.0001	0.024	0	-1.229	0.998
Prior Article Slant	Contribution	-0.056	0.208	0	-0.605	0.624
Prior Article Length	Contribution	4,128.249	3,757.940	3173	1	1,963,441
Prior Refs	Contribution	34.319	60.839	7	0	1,636
For Ideological Shift Analysis:						
Contributor Yearly Bias	Contributor-Year	0.004	0.028	0	0	0.819
Opposite-Slant Article Edits Fraction	Contributor-Year	0.006	0.051	0	0	1
Same-Slant Article Edits Fraction	Contributor-Year	0.069	0.207	0	0	1
Revision War Edits Fraction	Contributor-Year	0.014	0.056	0	0	1
Number of Edits	Contributor-Year	12.314	45.036	2	1	949
For Composition Shift Analysis:						
Stay	Contributor	0.030	0.170	0	0	1
Opposite-Slant Article Edits Fraction	Contributor	0.004	0.053	0	0	1
Same-Slant Article Edits Fraction	Contributor	0.075	0.251	0	0	1
Revision War Edits Fraction	Contributor	0.008	0.056	0	0	0.944
Starting Number of Edits	Contributor	0.321	0.693	0	0	9.568

Notes: Number of observations in this table is: 9,487,164 for contribution level; 381,099 for contributor-year level where the contributor made at least one edit in the year before; and 2,482,063 for contributor level where the contributor joined Wikipedia before 2010.

I.5. Analyzing Contributor Slant

Motivated by the pattern in the raw data, this section quantifies how contributors' contributions change with their edits and analyze whether (and how) their editing experience affects their slant decline.

[21] For contributor–year-level variables, observations include only contributor–year combinations where the contributor made at least one edit in the previous year, because the independent variables for the composition shift analysis are calculated based on the past year's edits.

I.5.1. Contributors' Participation Pattern on Wikipedia

We first investigate the type of content contributors interact with on Wikipedia. For every edit in our sample, we estimate the following regression model:

$$Contributor\ Slant_j = \alpha_0 + \alpha_1 Prior\ Article\ Slant_{ijt} + X_{ijt}B + \sigma_i + \eta_t + \varepsilon_{ijt}. \qquad (1)$$

In this baseline specification, we set the contributor slant to a fixed value, even though we can observe the same contributor multiple times. Coefficient α_1 identifies whether the average contribution follows EC or non-EC, as noted earlier. To address concerns about unobservable factors influencing the choice, we include X_{ijt}, a vector of the article's characteristics and control variables after article *i* is edited by contributor *j* at time *t*, and σ_i, an article-fixed effect, to control for any fixed differences among articles, and η_t, a year fixed effect, to control for any common media/macroeconomic shocks or Wikipedia policy changes that may differentially affect articles from different years. We note that the key exogenous variable is measured with considerable noise, which can induce attenuation bias in the estimate.[22] Hence, we view the result, at best, as an underestimate. In an alternative approach that mitigates the noise, we create two categorical variables. On the basis of *Contributor Slant*, we create *Contributor Category*, which takes the value of -1, 0, or 1, representing contributors with a slant two standard deviations below the mean, in between, and above the mean, respectively. *Prior Article Category* is the categorical version of *Prior Article Slant*. We use *Contributor Category* as the dependent variable, with *Prior Article Category* as the explanatory variable, to estimate standard models for categorical choice (see Appendix Table A1).

[22] This is a standard concern with poorly measured exogenous variables. See, for example, Draper and Smith, 1998, p. 19.

In Table 3, we report the estimation results of Equation (1) using ordinary least squares (OLS) regressions. Models (1) to (3) use *Contributor Slant* as the dependent variable. Model (1) includes only *Prior Article Slant* as the explanatory variable. Model (2) adds the control variables *Log (Prior Article Length)* and *Log(Prior Refs).* Model (3) replicates Equation (1), with article and year fixed effects included. The coefficients on *Prior Article Slant* are negative and significant in all three models.[23] This finding indicates that an increase in the article's slant is associated with a decrease in the slant of its next contributor. That is, when the article is more Republican leaning, it tends to attract a more Democrat-leaning user as its next contributor. This pattern is consistent with the non-EC behavior.

Table 3: OLS Regressions on the Relationship between Contributor Slant and Prior Article Slant

Model	(1)	(2)	(3)
Dependent Variable	Contributor Slant	Contributor Slant	Contributor Slant
Prior Article Slant	-0.0087*** [0.0001]	-0.0086*** [0.0001]	-0.0189*** [0.0004]
Log(Prior Article Length)		0.0006*** [0.0000]	0.0009*** [0.0001]
Log(Prior Refs)		-0.0004*** [0.0000]	-0.0010*** [0.0001]
Observations	9,487,164	9,487,164	9,487,164
Adjusted R-squared	0.006	0.006	0.007
Year FE	No	No	Yes
Article FE	No	No	Yes
Number of Articles	65,046	65,046	65,046

Notes: The unit of analysis is each edit in our analysis sample. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

[23] As a robustness check, we added the slant of the last edit on the article as a control variable and clustered the standard erorr at the article level. All results continue to hold.

I.5.2. Ideological Shift: How Does Editing Experience Change the Contributions from Contributors?

Why do biased contributors become moderate? What effect does opposite content have? We estimate the following:

$$Contributor\ Yearly\ Bias_{jt} = \beta_0 + \beta_1 Number\ of\ Edits_j + \beta_2 Opposite\ Slant\ Article\ Edits\ Fraction_{j,t-1} + \beta_3 Same\ Slant\ Article\ Edits\ Fraction_{j,t-1} + \beta_4 Revision\ War\ Edits\ Fraction_{j,t-1} + \mu_j + \epsilon_{jt}. \quad (2)$$

The unit of analysis is contributor–year, which allows us to observe how each contributor's ideology evolves over time. Observations in the sample include contributor–year combinations where the contributor made at least one edit in the previous year as the independent variables are calculated with reference to the past year's edits. The dependent variable is *Contributor Yearly Bias*, which captures the average bias of a contributor's edits in a given year. The contributor's yearly total number of edits, *Number of Edits*, controls for how active the contributor was in the past year.

To capture a contributor's experience in interacting with different types of content, we count the contributor's number of edits in a given time period targeting extreme opposite-slant articles, divided by the contributor's total number of edits during that time, as *Opposite-Slant Article Edits Fraction*, and the proportion of a contributor's edits targeting extreme same-slant articles, as *Same-Slant Article Edits Fraction*. Moreover, to investigate how pushbacks may shape contributors, we measure a form of pushback called the *revision war*, where a contributor's edit is immediately reverted by another contributor and the original contributor edits back the same contribution. We count the edits that a contributor makes during such revision wars, divided by the contributor's total number of edits, as *Revision War Edits Fraction*. In this regression model, we include the lagged *Opposite-Slant Article Edits Fraction*, *Same-Slant Article Edits Fraction*, and *Revision War Edits Fraction* to test how a contributor's experience in the past year with different types of extreme content and pushbacks affects his or her likelihood of adding more bias to existing content. μ_j is a contributor fixed effect to control for any fixed differences among contributors.

Table 4 reports the regression results. In Model (1), we observe that although encountering extreme content with the same slant reinforces the contributor's own ideology, interacting with opposite-slant extreme content causes the contributor's slant to become moderate. We do not find a significant effect from receiving pushbacks on the contributor's average bias. When we use an alternative dependent variable, *Contributor Yearly Maximum Bias*, in Model (2), computed the same as *Contributor Yearly Bias* but taking the contributor's maximum bias instead of average bias in the year, we find a significant negative effect of receiving pushbacks on the contributor's maximum bias in that year.[24] In summary, encountering extreme content of the opposite slant (rather than the same slant) or receiving pushback from other contributors reduces the contributor's own bias to some extent.

Table 4: OLS Regressions on Contributor Slant and Content the Contributor Interact With

Model	(1)	(2)
Dependent Variable	Contributor Yearly Bias	Contributor Yearly Maximum Bias
Log(Number of Edits)	0.0003*** [0.0001]	-0.0001 [0.0003]
Opposite-Slant Article Edits Fraction	-0.0536*** [0.0042]	-0.1599*** [0.0081]
Same-Slant Article Edits Fraction	0.0007* [0.0004]	0.0040*** [0.0012]

[24] We also added year fixed effects and clustered the standard erorr at the contributor level as a robustness check. All results continute to hold.

Revision War Edits Fraction	0.0005 [0.0013]	-0.0084* [0.0048]
Observations	381,099	381,099
R-squared	0.0003	0.0060
Year FE	Yes	Yes
Contributor FE	Yes	Yes

Notes: The unit of analysis is each contributor-year. Observations include contributor-years where the contributor made at least one edit in the year before. *Contributor Yearly Bias* is the average absolute value of the slants, i.e., the biases, of a contributor's edits that year. *Contributor Yearly Maximum Bias* is the maximum bias of a contributor's edits that year. A *Revision war* is defined as a contributor's edit being reverted immediately by another contributor, and then is immediately followed by the original contributor editing back the same contribution, as a "fight back." The "fractions" in this table are lagged by one year; for example, *Revision War Edits Fraction* in this table equals the number of edits that the contributor made in the past year during such pushback situations, divided by the contributor's total number of edits in the past year. Robust standard errors in brackets, clustered at the contributor level. *significant at 10%; ** significant at 5%; *** significant at 1%.

Interpreting the above coefficients is difficult, so we use a Markov matrix to illustrate how the slant composition of contributors evolves. This matrix, reported in Table 5, is constructed as follows. First, we divide the time that a contributor has been on Wikipedia in half. Then, we divide the direction of this contributor's edits by attaching values (-1, 0, 1) to negative, zero, and positive slant edits. On the basis of the sum of these values for the first and second halves of the contributor's activity, we categorize the contributor as Democrat, Neutral, or Republican. If the sum of all edits in one half is negative (positive), then the contributor is a Democrat (Republican), and if the sum of all edits in this half is zero, then the contributor is neutral. We perform this step for each half of every contributor's activity on Wikipedia and accumulate them to get the overall transition probabilities in the entire community. For the Democratic- and Republican-leaning contributors in the first half, there is more than a 70% chance that they will move to neutral in the second half of their contribution life span.

Table 5: Transition Matrix of Contributor Slant Change in Wikipedia

		First half of activity		
		Democratic Type	Neutral	Republican Type
Second half of activity	Democratic Type	0.1407	0.0328	0.1145
	Neutral	0.7451	0.9333	0.7416
	Republican Type	0.1142	0.0339	0.1439

Notes: The sample is constructed by dividing every contributor's time in half. Then divide the direction of his or her edits, i.e. attach values (-1, 0, 1) to negative, 0, positive slant edits. Sum up the edits' values for the first half and the second half of his or her activity. If the sum of all edits in this half is negative, the contributor is a Democrat Type in this half. If the sum of all edits in this half is zero, the contributor is Neutral in this half. If the sum of all edits in this half is positive, the contributor is Republican Type in this half.

Although the community of participants has a tendency of moving toward neutral, Table 5 does not provide any sense of whether this change is more or less pronounced than the composition shift. We first provide evidence for a causal explanation. Then we characterize the composition shift and compare the two shifts.

I.5.3. Mass Edits: Causal Evidence

The ideal design to establish causality is to employ some exogenous shocks and observe how contributors’ slants change before and after these shocks and compare these actions with those that did not receive shocks. We operationalize this idea with the special circumstances of *mass edits*, where an article attracts an unusually high number of contributions in one day due to a sudden social event or breaking news about the topic. On such occasions, the article usually receives a large volume of searches online. However, social events or breaking news is unpredictable, and so are the mass edits.

We define mass edits using online search volumes from the Google Trends website.[25] Specifically, we use the article's title as the search keyword(s) in Google Trends, collect the global daily Google Search Index (GSI) for a two-week window around the potential "mass edit" date, and compare whether the average GSI of the three-day window around the search date (days -1, 0, and 1) is greater than the three-day average GSI before the mass edit date (days -4, -3, and -2) and the three-day average GSI after the mass edit day (days 2, 3, and 4). We then aggregate all the contributions to the article–date level and define an article as experiencing a "mass edit event" if it meets the following conditions: 1) the article receives more than 10 contributions on that date and 2) its title has an abnormal search peak in GSI during the three-day window, as described earlier. We also combine consecutive mass events that happened to the same article on multiple consecutive days as a single mass edit event. Finally, we focus on the top mass edit events whose number of edits is above the 99th percentile of all events as they best represent such shocks. In this way, we identify the top 35 mass edit events.

To estimate the effect of experiencing mass edits on a contributor's later slant, we use a treatment/control approach. We consider the contributors who are exposed to a given mass edit the *Treated* contributors. We use a propensity score matching method to construct a control group of contributors. Specifically, for a given mass edit event, we identify all the contributors who made at least one edit during the event. Then, for each of the treated contributors, we find another contributor who did not edit during the event to create a "matched pair" in our data set. We perform matching based on their previous slant and editing behavior before the mass edit event. Once a contributor is matched, we exclude him or her from future matching, and we repeat the process for

[25] For descriptions of the Google Trends website and Google Search Index, see https://support.google.com/trends/?hl=en#topic=6248052, accessed May 2018.

the next mass edit event. Altogether, we identify 5,148 treated contributors who experienced mass edits and another 5,148 control contributors as their matched pairs.

First, we examine whether mass edits produce more biased content, which may lead to contributor slant change. We use a t-test to examine whether the articles being edited during mass edit events are more biased compared with those under normal edits. The results show that, on average, articles ("article versions") after mass edits are more biased than those experiencing normal edits, that is, the average absolute article slant for articles after mass edits (0.1346) is significantly greater than the average absolute article slant after normal edits (0.0839), with p = 0.000. Indeed, mass edit events are more likely to produce biased content.

Next, we use OLS regressions to examine how the treated contributor's average bias changes compared with the control contributors due to the mass edits. Table 6 reports the results. The treated contributors' average biases after experiencing the mass edit event are significantly lower compared with those of their pairs who were not exposed to mass edits.

Table 6: OLS Regressions on How Contributor Slant Changes during Mass Edits

Model	(1)
Dependent Variable	Contributor Bias
Treated	-0.0015*** [0.0005]
Observations	10,296
Adjusted R-squared	0.001

Notes: Observations are the contributors who experienced mass edit events and their matched pairs. *Contributor Bias* is the contributor's average absolute value of their edits after the mass edit event. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

Why do the biases of contributors decrease after experiencing mass edits? In our proposed mechanism, encountering extreme opposite-slant content reduces contributors' bias. Contributors may be more frequently exposed to extreme opposite-slant content during mass edits, which leads to their ideological shift. To test this mechanism, we perform a simple t-test mean comparison to compare the average article slant during mass edits vs. normal edits. The results show that the article versions after each mass edit are significantly *more biased* than those after each normal edit ($p < 0.0000$). This finding means that contributors are exposed to more extreme content during mass edits.

Next, we define a "flip" of an article's slant when the slant changes from extremely left/right leaning (i.e., more than two standard deviations away left/right from neutral) to extremely right/left leaning. We estimate the likelihood that an article has at least one slant flip on mass edit days and compare it to normal edit days. The logit regressions are shown in Table 7. Aggregated to the article–date level, *Flip* is equal to 1 if an article has at least one flip on a given day and 0 if the article has no slant flip on that day. After converting the estimated coefficients, an article is 11.8% more likely to experience slant flips during mass edits than during normal edits.

Table 7: Logit Regressions between Article Slant Flips and Mass Edit Events

Model	(1)	(2)
Dependent Variable	Flip Dummy	Flip Dummy
Mass Edits Dummy	1.4199*** [0.1151]	2.0178*** [0.1161]
Prior Article Slant		-0.6371*** [0.0657]
Log(Prior Refs)		-0.0868*** [0.0117]
Log(Prior Article Length)		-0.4402*** [0.0079]
Observations	5,804,714	5,804,714
Pseudo R-squared	0.001	0.038

Notes: Observations in this panel are at the article-date level. The dependent variable *Flip Dummy* equals 1 if the article experiences at least one slant flip during that day.

Mass Edits Dummy represents whether the article is receiving mass edits on the given day. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

The analyses provide additional evidence to support our proposed mechanism. Contributors encountering mass edits are more likely to be exposed to content of both extremes. Because the impact of the opposite-slant extreme content on a contributor's slant is greater than that of the same-slant extreme content (Table 4), bias decreases more in the contributors who encounter mass edits compared with those who do not. The exogenous nature of mass edit events allows us to obtain a causal inference of such effects.

I.5.4. Composition Shift: Why Do Extreme Contributors Leave Over Time?

To examine each contributor's exit decision, we employ the following Logit regression to examine each contributor's likelihood of staying/exiting:

$$Stay\_Dummy_j = \theta_0 + \theta_1 Starting\ Number\ of\ Edits_j + \theta_2 OppositeSlant\ Article\ Edits\ Fraction_j + \theta_3 SameSlant\ Article\ Edits\ Fraction_j + \theta_4 Revision\ War\ Edits\ Fraction_j + Vintage\ Dummies_j + \omega_j, \quad (3)$$

where the unit of analysis is each contributor. We include a similar set of explanatory variables as shown in Table 4. The contributor's number of edits in the first two years, *Starting Number of Edits*, controls for how active the contributor was when joining Wikipedia. We also control for the contributor's year of entry to capture the vintage effect.

Table 8 presents the results. Model (1) includes all contributors who joined before 2010, whereas Model (2) includes only the core contributors. The coefficients for *Opposite-Slant Article Edits Fraction* and *Revision War Edits Fraction* are negative and statistically significant.

Calculating the average marginal effects shows that experiencing the opposite-slant extreme content or revision wars reduces a contributor's likelihood of staying on Wikipedia by 2.7% and 25.2%, respectively, compared with a contributor who encounters no opposite-slant content or revision wars. The effect of interacting with same-slant content, shown by the coefficient of *Same-Slant Article Edits Fraction*, goes in the other direction, which increases the likelihood of a contributor staying. The results in Tables 4, 6, and 8 show that encountering more extreme opposite-slant content or pushbacks from others leads contributors to become either more moderate or more likely to leave.

Table 8: Logit Regressions on Contributors' Likelihood of Staying on Wikipedia

Model	(1)	(2)
Sample	All Contributors	Core Contributors Only
Dependent Variable	Stay	Stay
Log(Starting Number of Edits)	1.6190*** [0.0049]	0.6539*** [0.0051]
Opposite-Slant Article Edits Fraction	-1.2807*** [0.1258]	-0.6001*** [0.0796]
Same-Slant Article Edits Fraction	0.0688*** [0.0156]	0.0592** [0.0268]
Revision War Edits Fraction	-11.7873*** [0.2883]	-9.3112*** [0.1258]
Observations	2,482,063	337,638
Pseudo R-squared	0.316	0.127
Vintage Dummies	Yes	Yes

Notes: The unit of analysis is each contributor. Observations include contributors who joined before 2010. Model (2) includes only core contributors, i.e., contributors who made at least 3 edits in the sample. The dependent variable *Stay* equals 1 if the contributor made at least 1 contribution in 2010 or 2011. *Starting Number of Edits* is the total number of contributions that a contributor made in the first two years after he or she joined Wikipedia. *Vintage* is the year in which the contributor joins Wikipedia, or the vintage that the contributor belongs to. The "fractions" are defined the same as in Table 4, but for contributor's edits in all time instead of the year before. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

I.5.5. Which Effect Contributes More to Changes in the Overall Bias?

We compare the size of each effect using various five-year subsamples to identify which effect contributes more to overall trends. Below, we take the five-year period of 2004–2009 to illustrate the simulation, and we repeat this process for all other five-year periods.

The goal is to calculate the proportion of the change in overall bias that is due to the ideological shift and that due to the composition shift. To begin, we compute the average bias—the absolute value of slant—of any contributor who contributed at least once in the first and last years. We denote the former crowd of contributors as the *starting crowd* and the latter crowd of contributors as the *ending crowd*. The average bias of the starting crowd in 2004 is 0.00606, whereas the average bias of the ending crowd in 2009 is 0.00421, a decline of 30.5%.

Next, we identify the composition of each crowd. In the starting crowd, 2,866 (4.34%) contributors are the staying contributors who still edit after five years, with an average bias of 0.00539, and 63,124 (95.66%) are the leaving contributors who did not edit after five years and whose average bias in the first year is 0.00609. The average bias of the starting crowd is 0.00606 = 0.00539 × 4.34% + 0.00609 × 95.66%. In the ending crowd, the 2,866 staying contributors have an average bias of 0.00346. The remaining 522,985 (99.45%) are new contributors who joined during the five-year period, whose average bias in 2009 is 0.00421. Thus, we compute the proportion of the bias decline due to each effect.

To consider the importance of the composition effect, we simulate the possible scenario if only the composition effect shaped the outcome. If some contributors leave after five years but the remaining contributors have the same slant as in the beginning, then the ending crowd's bias would be equal to the average of the staying contributors' beginning bias (0.00539) and the new contributors' bias (0.00421), yielding 0.00539 × 0.55% + 0.00421 × 99.45% = 0.00422.

Accordingly, the five-year bias decline from the composition effect will be 0.00606 - 0.00422 = 0.00184.

For comparison, we simulate the possible scenario if only the ideological shift shaped the outcome. To estimate the effect of ideological shift, we observe all contributors who are still edit after five years but their slant changes over this time. This is the crowd whose ideologies shift. If we assume that such ideological shift happened to all contributors in the starting crowd instead of just those who remain active, then the ending crowd's average bias would be equal to the average of the staying contributor's decreased bias (0.00346) and the new contributors' bias (0.00421), yielding 0.00346 × 11.2% + 0.00421 × 88.8% = 0.00413.[26] Accordingly, the five-year bias decline from the ideological shift only will be 0.00608 - 0.00413 = 0.00195.

The ideological effect is not common enough to have a large effect on the overall results. It involves as few as only 4.34% of the starting crowd in the above calculation. With an actual crowd bias decline of 0.00185 in the raw data, a simple equation solving for the proportion of each effect yields the estimate that roughly 88.4% of the change is due to the composition effect and the remaining 11.6% is due to the ideological shift effect. Although this is a simple back-of-the-envelope calculation, it provides an indicator that the decline in contributor slant is largely due to the composition effect.

We repeat the above simulation for all five-year periods in our sample and report the effect sizes in Table 9. Panel A uses the same sample as in the main analysis, which excludes the first edits of each article that creates the article, and Panel B uses the full sample containing all the raw edits. Apart from the first five-year period, both sets of results show a consistent pattern: composition shift accounts for 80.25%–90.98% of the effect, whereas ideology shift accounts for

[26] 11.2% = (2,866 + 63,124) / (2,866 + 63,124 + 522,985).

the remaining proportion of the effect ranging from 9.02% to 19.75%. The effect sizes are robust across different five-year periods except for the first five years (i.e., 2001 to 2006). A possible reason is that in 2001, when Wikipedia was first founded, the earlier contributors contain a larger portion of extreme participants than later years, as is also shown by the percentages in Table 1. The total number of contributors that year is also only 0.03% of the full sample. As this crowd has a large portion of extreme contributors to begin with and with the extreme contributors becoming neutral over time, a greater ideology shift is observed in the first five-year period compared with later years. The later subsamples have more participants and examine Wikipedia in its most "developed" state, where the participants are more familiar with all the norms and rules.

Table 9: Simulation Results Comparing Ideology Shift and Composition Shift over Time

Panel A – Percentages of the Effect Sizes over Time, Using Analysis Sample

Five-Year Period	Composition Shift	Ideology Shift
2006-2011	83.56%	16.44%
2005-2010	83.95%	16.05%
2004-2009	88.36%	11.64%
2003-2008	90.98%	9.02%
2002-2007	89.38%	10.62%
2001-2006	48.03%	51.97%

Panel B – Percentages of the Effect Sizes over Time, Using Full Sample

Five-Year Period	Composition Shift	Ideology Shift
2006-2011	81.72%	18.28%
2005-2010	81.62%	18.38%
2004-2009	87.70%	12.30%
2003-2008	89.07%	10.03%
2002-2007	80.25%	19.75%
2001-2006	50.58%	49.42%

Notes: Percentages reported are from the simulation in Section 5.5 comparing the effect sizes of Composition Shift and Ideology Shift. Panel A uses the same sample as in the main analysis, which excludes the first edit of each article that creates the article. Panel B uses the full sample containing all the raw edits.

I.6. Discussion

I.6.1. Rate of Slant Change: How Long Will It Take for Contributors to Become Neutral?

We estimate how long it takes for a contributor's slant to gradually become neutral if this tendency continues. In this simulation, we observe the slant at an aggregated level, and we simulate the slant change of Wikipedians as a crowd, regardless of which part of the change is due to the ideological or composition shift.

We use a Markov chain process to simulate the evolution. Although a contributor's slant exhibits a long-term trend over the years, it frequently fluctuates, and this fluctuation should be accounted for. We divide the slant into different bins and investigate how a contributor's slant changes from one bin to another. *Contributor Yearly Slant* is divided into seven bins and divided by the ± 0.5, ± 1.5, and ± 2.5 standard deviation intervals. The middle bin represents a neutral slant, whereas the first and last bins represent the extreme slants. We then compute a transition matrix for contributor slant based on our empirical data. For each year, we compute the proportion of contributors whose yearly slant moves from one slant bin to another and fill the probabilities in the transition matrix for this year. Averaging the transition matrices across all years gives us the final transition matrix we use in our simulation (reported in Table 10).

Table 10: Transition Matrix of Contributor Slant Change over Time

		Start Slant						
		bin1 [-1.229, -0.059)	bin2 [-0.059, -0.035)	bin3 [-0.035, -0.012)	bin4 [-0.012, 0.012)	bin5 [0.012, 0.035)	bin6 [0.035, 0.059)	bin7 [0.059, 1.000)
End Slant	bin1 [-1.229, -0.059)	0.8298	0.0139	0.0024	0.0011	0.0013	0.0008	0.0015
	bin2 [-0.059, -0.035)	0.0717	0.7242	0.0044	0.0020	0.0103	0.0019	0.0007
	bin3 [-0.035, -0.012)	0.0591	0.1745	0.7438	0.0055	0.0040	0.0149	0.0029
	bin4 [-0.012, 0.012)	0.0323	0.0713	0.2286	0.9795	0.2089	0.0531	0.0277
	bin5 [0.012, 0.035)	0.0036	0.0128	0.0177	0.0060	0.7545	0.1867	0.0624
	bin6 [0.035, 0.059)	0.0008	0.0014	0.0015	0.0033	0.0052	0.7222	0.0757
	bin7 [0.059, 1.000)	0.0028	0.0019	0.0018	0.0025	0.0158	0.0203	0.8291

Notes: *Contributor Yearly Slant* is split by the ± 0.5, ± 1.5, and ± 2.5 standard deviations intervals. The middle bin represents neutral slant; the first/last bin represents extreme slant.

In this transition matrix, the rows denote the starting bins and the columns denote the ending slant. Bin 4 represents a neutral slant, defined as a slant index ranging from -0.5 to 0.5 standard deviations away from the mean. We find that: (1) the probabilities on the diagonal are large. As expected, contributors tend to have a high chance of staying near their original slant; and (2) the farther the end bins are from the start bins, the smaller the probabilities. These findings indicate that the contributor slant change is a gradual and cumulative process, and it is not likely that the contributor's slant would suddenly jump from one extreme to another.

Next, we use the transition matrix to simulate the contributor slant change process over time. We compute the time it takes for a contributor to have a greater than 50% probability of moving to neutral (see Table 11). As expected, the length of time depends on the contributor's original slant: Extremely slanted contributors spend a longer time moving to neutral than slightly slanted contributors. Surprisingly, we find that, on average, it takes one more year for Republicans to become neutral than for Democrats.

Table 11: Time Needed for a Contributor to Have > 50% Probability of Moving to Neutral Slant

Starting Contributor Slant	**Number of Years**
Extremely Democratic	10
Democratic	6
Slightly Democratic	3
Neutral	0
Slightly Republican	4
Republican	7
Extremely Republican	11

Notes: Number of years calculated based on the Markov Chain Process. *Neutral* state includes contributor slant 0.5 standard deviation away from 0. *Slightly Democratic (Republican)* state includes contributor slant between 0.5 and 1.5 standard deviations below (above) 0. *Democratic (Republican)* state includes contributor slant between 1.5 and 2.5 standard deviations below (above) 0. *Extremely Democratic (Republican)* state includes contributor slant more than 2.5 standard deviations below (above) 0. On average, after about 30 years, the probabilities in all articles' end state reach stationary distribution, with the probability of contributor slant moving to *Neutral* being 87.4%.

We test for possible reasons why Republican contributors tend toward a neutral slant more slowly than Democratic contributors. First, do Republican contributors display more EC behavior than Democratic contributors? The regression results of Equation (1) using the two groups separately do not support this explanation. Republican contributors show less EC behavior than Democratic contributors.

Second, Republican contributors may choose to edit less extreme articles compared with Democratic contributors, and so they are less influenced during their interaction with the online content. However, no statistically significant difference was found between the level of content extremeness for the articles edited by Republicans or Democrats. The distributions contain similar bias and variance.

A third possible reason may stem from the contributors' number of edits. Republican contributors make fewer edits than Democrats, so their experience has less of an effect on the overall tendency and may differ in some way. Summary statistics provide evidence for this explanation. In our sample, the total number of edits from Democratic contributors is approximately 1.5 times that of Republican contributors.

I.6.2. Is the Measure of Contributor Slant Representative of Ideologies?

One may be concerned about whether the measure of slant in Wikipedia is representative of contributors' real-world political ideologies. In addition, a neutral article in our sample can either be interpreted as having no slanted words at all or as having equal numbers of very slanted words. These concerns may lead to questioning the external validity of the slant measure.

To address these concerns, we use an alternative measure of the slant and bias. We match the voting data from the 2004 presidential election to locations affiliated with contributors' IP addresses.[27] We restrict our sample to contributors who are not logged in when editing the articles because Wikipedia only reveals IP addresses for contributors without user IDs. We also drop contributors with IP addresses located outside the U.S.. We then test the relationship between the voting record and *Prior Article Slant* using OLS regressions. Note that this analysis analyzes the behavior of a different population of contributors than the contributors we have examined thus far.[28] This regression is valid under the assumption that a contributor has, on average, the political preferences of the region in which he or she lives.

[27] The data on the geolocation of IP comes from MaxMind. We match up the county records.

[28] The identities of contributors are known after they register and when they edit after logging on. An anonymous edit comes from either an unregistered contributor or from an editor who chose not to log on before editing. Hence, the samples can possibly include some of the same contributors, but identifying the fraction is impossible.

Appendix Table A2 presents the results. *RepPerc* denotes the percentage of Republican votes in the contributor's county. As we use positive values in the slant index to indicate Republican-leaning ideologies for Wikipedia users and articles, the negative and statistically significant coefficient of *Prior Article Slant* suggests that a contributor from a county with a higher percentage of Republican votes tends to target a Democratic-leaning article when he or she contributes on Wikipedia. The results show a non-EC pattern in the contributing process and are qualitatively similar to the prior estimates. This finding also supports the notion that the measure of contributors' slant reflects the contributors' real political ideologies.

We also collect talk pages for articles, which are used by contributors to discuss edits and achieve consensus. The total size of an article's talk pages has a correlation of 0.22 with the average bias of the article over time, suggesting that our bias measure does capture how contested an article is.

I.6.3. What Else Could Be Driving the Non-EC Behavior?

The effect of non-EC in contributors' voluntary editing behavior indicates that contributors are more likely to edit articles with the opposite slant. This can also be due to the revision war among contributors, which may have little to do with the article's slant. We address this concern by including only the initial edits of every contributor when they revise an article for the first time. This rules out revision wars or any possible correcting behavior later in the edits.

Appendix Table A3 shows that the signs and statistical significance of the estimated coefficients do not change, and the magnitude of the coefficients becomes even larger, indicating an even stronger non-EC effect than that when investigating all edits. The results further strengthen the robustness of the effect.

I.7. Conclusions

Wikipedia has a long record of bringing opposing opinions into the same conversation. Over time, Wikipedia has become less biased. Our study finds that this change is partly due to an ideological shift, where contributors contribute less slanted content over time. It is also due to a composition shift, where extremely biased contributors contribute less content over time. Contributors interact with opposite viewpoints and do so more frequently than participating in ECs. Extreme contributors either become more moderate or tend to leave after interacting with an opposite-slant content or encountering pushbacks from others.

This study offers an approach for identifying the mechanisms contributing to (un)segregated conversations. It identifies the factors that alter the composition of the crowd and cause a contributor's viewpoint to evolve. Nothing in this approach presumes the results; the approach can flexibly measure contributions to (un)segregated conversations in a variety of settings.

The findings inform open questions for two important managerial issues. Wikipedia's stewards, the Wikimedia Foundation, face an important question about how to encourage the emergence of an NPOV. Our findings suggest that retaining existing contributors with moderate opinion is, by far, the most important factor in maintaining an NPOV on its articles. Wikipedia should continue to expose all extremes to the opposite opinion as it tends to arise from normal practice, which leads to the exit of the most extreme contributors and moderation of views among those who stay. This approach will work as long as the entry of new contributors continues to draw on a diverse set of opinions as it has in the past.

For webmasters of crowds, we draw a related lesson from Wikipedia's experience. If a webmaster aspires to draw on multiple opinions and achieve near-neutrality in the content produced by their online communities, then the experience at Wikipedia does not suggest a passive approach to managing contested knowledge. Simply maximizing participation, regardless of the opinion, is also a mistake. Webmasters must articulate their aspirations for an NPOV and insist that contributors also aspire to that goal while recruiting a diversity of opinion. If they are successful at recruitment, then actively discouraging participation from those who maintain extreme points of view is reasonable. Indeed, our findings suggest, cautiously, that one way to achieve moderate outcome is to encourage extremists to leave. We stress the cautious nature of the advice because Wikipedia samples from wide and multiple points of view in spite of changes in the composition and that a wide sample should not be taken for granted in other settings.

The findings also raise a subtle question: How does Wikipedia transform controversial topics into arguments that include many points of view and sustain the community over time? We speculate that Wikipedia's success in this regard arises from the institutions that help overcome the challenges affiliated with aggregating contested knowledge. First, the aspiration of achieving an NPOV directs attention to the principle that no side can claim exclusive rights to determine the answer. Second, norms allow every contributor to add another paragraph if it reduces tension by giving voice to dissent. Reducing disputes this way costs little: miniscule storage and transmission costs reduce the cost of listing another view on a web page. Our results also suggest that the conflict resolution mechanisms and the mix of informal and formal norms at Wikipedia play an essential role in encouraging a community that works toward an NPOV. This finding is consistent with theories suggesting that articles go through a lifecycle and settle into a consensus, which contributors subsequently "defend" (see, e.g., Kane et al. 2014). We also note a new open question: Although our findings suggest that Wikipedia's mechanisms work as desired, our findings raise questions about which specific norms, other than the declaration of principles, also contribute.

These findings also raise concerns on the platform design literature. We speculate that some simple design differences may have profound consequences for (un)segregating conversations. For example, on Facebook, an algorithm selects content for users, and its design increases the chance that the participants read and write content only in a community of like-minded people. By contrast, Wikipedia contributors can select to examine whichever content they desire and add and remove material or refine the content in myriad ways. Contributors on Facebook/Twitter can only add additional content on top of what is already there. Allowing the removal or editing of anyone's contributions can change how the reader and writer choose to direct the conversations, resulting in contributions from different points of view. Some platforms also aggregate contributions in ways that shape the prevalence of segregation. For example, on Yelp (e.g., rating restaurants) or Rotten Tomatoes (e.g., rating movies), additional materials can be added without limit. These platforms provide a numerical summary that can direct conversations between readers and reviewers. Our results prompt questions about whether a numerical summary motivates others with views that differ from the summary or attracts more reviews from those who agree with it, and how such a process makes the summary more valuable.

CHPATER II:

Trust and Disintermediation:

Evidence from an Online Freelancing Marketplace

II.1. Introduction

Intermediaries are everywhere in our economy: brokers in the finance and insurance industries, headhunters in the labor market, distributors in retail, housing agents in real estate, and online platforms in the information technology industry, just to name a few. In 2010, intermediaries contributed an estimated 34% of the US gross domestic product (Spulber 2011). Economists have long recognized the importance of intermediaries for providing matching and facilitating transactions (e.g., Spulber 1996a, 1996b, 1999, Parker and Van Alstyne 2005, Armstrong 2006, Rochet and Tirole 2006, Edelman and Wright 2015, Hagiu and Wright 2015). However, all intermediaries face the risk of disintermediation, in which two sides circumvent the intermediary to transact directly and avoid the intermediary's fees.

Disintermediation is prevalent. For example, the traditional role of book publishers as intermediaries was weakened when Amazon enabled authors to sell directly to readers through its self-publishing services. Li & Fung, a supply-chain management company that connects global retail brands with Chinese manufacturers, suffered ongoing decline in revenue as retailers disintermediated to work with manufacturers directly. A survey by ZBJ.com, the largest online freelance marketplace in China, indicates that approximately 90% of transactions are conducted outside the platform after clients and freelancers have been matched on its platform (Zhu et al. 2018). Hotels and airlines offer incentives to lure customers to book directly with them, thereby shrinking the revenue for online travel agencies. A few intermediaries have been unable to sustain their businesses as a result of disintermediation. For example, online platforms such as Homejoy went out of business because their consumers transacted with service providers outside the platforms.

Despite the importance of disintermediation with regard to firms' strategies and survival, there is scant literature on the issue, perhaps because of the difficulty of observing and measuring disintermediated transactions. In this paper, we study the relationship between trust and disintermediation, leveraging a randomized control trial (RCT) in an online platform. We utilize conversations recorded online to provide direct evidence of users' intentions to disintermediate. A number of studies have shown that building trust between two sides is crucial for platforms to facilitate effective matching among users (e.g., Resnick and Zeckhauser 2002). However, significant trust can reduce the perceived importance of platform services such as monitoring transactions, escrow payments, dispute settlements, and refunds for failed transactions (e.g., Edelman and Hu 2016). Users who trust one another feel less of a need for such services and are incentivized to take their transactions off the platform to avoid intermediary fees.

Our research setting is a large outsourcing platform that enables clients to find freelancers who satisfy the clients' job requirements. The platform then provides features through which the two sides contract, collaborate, create invoices, and pay, and it charges a per-transaction service fee that is approximately 10% of each transaction's value. In an RCT, the platform provider shows freelancers' satisfaction scores (SSs) to a random sample of clients. SSs are a newly developed measure of a freelancer's business reputation based on his or her complete work history on the platform. We find that the enhanced trust derived from seeing high SSs increases the likelihood for a high-quality freelancer to be hired. However, it also increases disintermediation between clients and freelancers with high SSs, as evidenced by significantly lower charges, fewer hours reported, and the stronger intention to disintermediate expressed in chat messages between them.

We conduct robustness checks to confirm that the reduction in hours and total charges is driven by disintermediation and not by other factors such as clients' selections of more efficient freelancers. Ultimately, increased disintermediation offsets the revenue gains from the increased hiring of high-quality freelancers. We also find that the tendency to disintermediate increases when users are geographically collocated, jobs are easily divisible, and clients themselves have high ratings.

Our study is related to the literature on trust within online platforms. Existing literature emphasizes the importance of building trust for successful business transactions among strangers online (e.g., Strader and Ramaswami 2002, Dellarocas, 2003, Pavlou and Dimoka 2006, Jin and Kato 2006, Cabral and Hortaçsu 2010, Cai et al. 2013, Moreno and Terwiesch 2014, Puranam and Vanneste, forthcoming). Although many factors may influence trust—such as an individual's past experiences transacting with the same person, escrow services by the platform, and certifications from trusted parties—in most studies, trust is reduced to the use of reputation systems (e.g., ter Huurne et al. 2017). Many experimental and observational studies provide evidence that reputation systems can effectively enhance trust (e.g., Ba and Pavlou 2002, Resnick and Zeckhauser 2002, Bohnet and Huck 2004, Pavlou and Dimoka 2006, Utz et al. 2009, Charness et al. 2011, Bolton et al. 2013). Such systems operate by mitigating the uncertainty involved in transactions, which stems from information asymmetry and potential opportunism between two transacting parties (e.g., Pavlou et al. 2007).

Further, a number of studies have identified the shortcomings of existing reputation systems in building trust. Research on eBay's feedback mechanism has shown that disappointed buyers often do not leave feedback (Nosko and Tadelis 2015), while buyers who have had a good experience are more likely to leave feedback (Masterov et al. 2015). Bolton et al. (2013) argue that reciprocity in providing feedback distorts reputation information, as fear of retaliation may deter users from truthful reporting. Even in a simultaneous-reveal system, in which reviews are not revealed until both parties have submitted their ratings, users may still be reluctant to provide negative feedback if they suspect it would discourage other parties from transacting with them (Luca 2017). These factors collectively result in reputation inflation. Ert et al. (2015) find that 97% of Airbnb ratings are between 4.5 and 5 stars; consequently, online reviews have no effect on the prices of Airbnb listings. Horton and Golden (2015) find similar patterns using data from Upwork.

Scholars have proposed various ways to develop better reputation measures in order to improve trust and transaction quality (e.g., Hui et al. 2014, Kapoor and Tucker 2017, Dai et al. 2018). However, our study suggests that under the disintermediation threat, attempts at enhancing the trust by providing a more accurate reputation system may actually harm the platform's ability to capture value.

Our study also improves our understanding of how different types of trust affects users' behavior. Bapna et al. (2017), building on Sobel (2005) and Cabral et al. (2014), categorize trust into intrinsic trust and instrumental trust. Intrinsic trust is motivated by the psychological benefits that an individual derives from being kind to others, while instrumental trust is backed by the option of rewarding and punishing a trustee in the future. Consistent with instrumental trust, we find evidence that even after clients and freelancers choose to disintermediate, they still prefer to initiate jobs on the platform so that they have the option to provide each other feedback or seek help from the platform.

Our study also adds to the small literature on disintermediation. Waldfogel (2012), Waldfogel and Reimers (2015), and Peukert and Reimers (2018) are three related studies that study the impact of digital disintermediation on product variety and quality in the music and book publishing industries. A few papers in the supply-chain management literature examine supplier encroachment as a means of circumventing intermediaries (e.g., Arya et al. 2007). Literature on brokerage also points out that brokers are more likely to be disintermediated when they do not control information (Rider and Samila 2019). However, none of these studies discuss how the degree of disintermediation changes with increased trust.

The remainder of the paper is organized as follows. Section 2 describes our empirical setting and design. Section 3 describes the data and variables. Sections 4 and 5 present empirical results and various robustness checks. Section 6 concludes by discussing limitations and managerial implications of this study.

II.2. Background and Empirical Design

The empirical context of our study is a major online outsourcing platform. A number of studies have examined the value of such platforms in online hiring (e.g., Agrawal et al. 2016, Stanton and Thomas 2016) and for conducting experiments (e.g., Horton et al. 2011). Jobs posted on the platform encompass a wide range of categories, such as Web, Mobile & Software Development, Design & Creative, Translation, Administrative Support, Accounting & Consulting, Writing, and Customer Service.

As soon as a client posts a job, a *job opening* is created, which typically includes a name, work description, requirements, and deadline. Any freelancer can submit a proposal to the client to bid for the job. Once the client selects a freelancer, the job is *filled* and a service contract (referred to as a *job assignment* hereafter) is created. A job assignment remains active until both parties agree to close it. Figure 9 illustrates the process flow of a typical job on the platform.

Figure 9. Job Process Flow

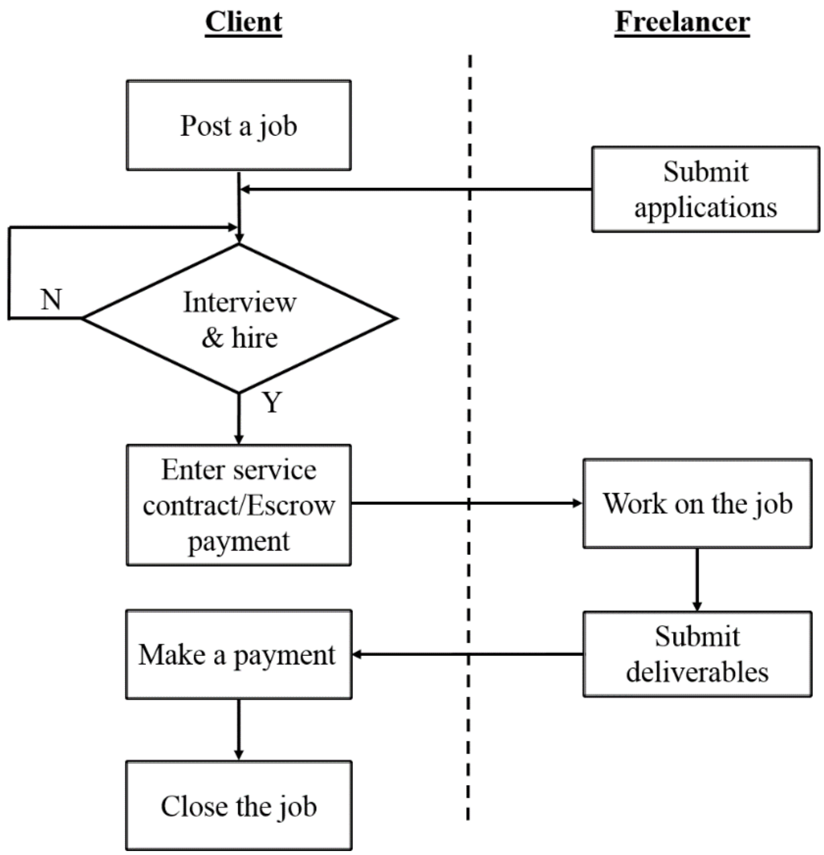

Clients can post either fixed-price or hourly jobs. The price for a fixed-price job is negotiated and determined between a client and a freelancer at the time of contracting; they can agree that the client will pay the total amount upon project completion or pay in stages according to agreed-upon milestones. For an hourly job, an hourly rate is decided at the time of contracting. Thereafter, the freelancer can begin working on the job and record working hours. After the freelancer makes a request for payment, the platform charges the client and holds the payment in escrow. The client has four days to review and dispute the amount. Once the dispute period ends, the escrow fund is released to the freelancer.

The platform charges freelancers a service fee of approximately 10% of the amount billed to the client. Disintermediation can take place in two ways. First, clients and freelancers can "chat" with each other on the platform at any time. Thus, they could agree to take jobs off the platform before initiating any projects to avoid the service fee. Second, durative transactions enable them to begin part of the job on the platform and then disintermediate for the remainder of the job to reduce the service fee. The latter approach to disintermediation enables clients and freelancers to leave each other reviews after they mark the job as complete. Our study focuses on the latter scenario.

Since its founding, the platform has used a basic five-star rating system to reflect user satisfaction. Star ratings are shown for both clients and freelancers on their user profiles, which reflect the average of all ratings received from completed jobs. This system experiences the shortcomings documented in the literature. First, the average rating from clients' past reviews does not take into account jobs for which no rating was given, nor does it allow for weight differentiation between older ratings and more recent ratings (e.g., Dai et al. 2018). Second, such a system may encourage reciprocity of positive reviews. Consequently, the average rating of freelancers on the platform is very high (above 4.5 out of 5). Thus, ratings do not accurately identify high- and low-quality freelancers. Finally, the five-star rating system rewards freelancers who have completed a large number of small and short-term projects and disfavors those who have worked on extensive and long-term projects.

For these reasons, the platform designed SSs as a new measure of freelancers' reputation, which represents a more complete picture of a freelancer's business. To avoid strategic manipulation by freelancers, the company does not disclose how the score is determined. However, the company does make it explicit that in addition to the ratings of past jobs, the SSs capture the following information that five-star ratings do not capture: 1) private feedback that clients have provided to the platform,[29] 2) disputes that freelancers have had with past clients, and 3) the number of times clients chose not to provide ratings. While clients can observe point 3) by browsing a freelancer's past work history, clients have no means of obtaining the first two information sets before the introduction of SSs. Research has shown that private ratings can help overcome concerns of retaliation and reciprocity and, thus, encourage clients to reveal truthful information (e.g., Luca 2017). Disputes frequently lead to project cancellation, in which case, the cancelled projects are not included in freelancers' work history, and neither clients nor freelancers leave reviews for each other. However, this information can be very valuable to clients for evaluating the risks associated with working with a given freelancer. Because of these advantages over five-star systems, a high SS can significantly boost a client's confidence that the freelancer can successfully complete the job.[30]

We leverage an RCT that the platform conducted from February 13 to March 10, 2015. It included a random sample (approximately 3%) of registered clients on the platform. The randomization is at the client level, and 50% of the sample clients were selected as the treatment group. When clients in the treatment group logged on to the website and browsed for freelancers, they were shown an SS in addition to the star rating on each freelancer's profile, while clients in the control group saw only the star ratings (Figure 10). The platform did not disclose the use of SSs to clients in the control group or to any freelancers. Moreover, freelancers were unable to observe their own SSs. Given the short duration of the experiment, we expect information leakage to be negligible. In addition, because freelancers were not notified of the introduction of SSs nor

[29] Private feedback is collected at the same time as public feedback when the job is closed, and both parties are informed that the private ratings are visible only to the platform.

[30] SSs are only implemented for freelancers, because providing better services to clients is a priority for this platform.

could they observe their own scores, we could disregard the possibility of signaling in their job applications. Figure 11 illustrates the distribution of SSs in the assignment sample. We collect data for all jobs initiated during the trial, even if a job is completed after the trial ends.

Figure 10. User Information Shown to the Treatment and Control Groups

Figure 11. Distribution of Freelancer Satisfaction Scores

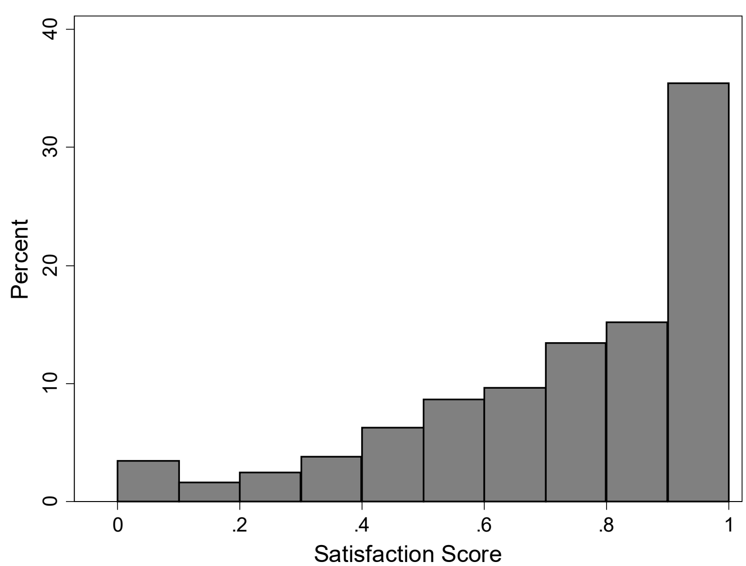

II.3. Data and Variables

We collect all job openings and assignments created by 24,732 clients from the treatment group and 24,458 clients from the control group during the trial. The analysis sample is at the job-assignment level, consisting of each assignment's outcome and the characteristics of the corresponding client and the hired freelancer. Jobs that were not filled or observations in which the freelancer's SS was not available are dropped, leaving a final sample of 33,561 job assignments.[31]

We employ two approaches to measure disintermediation. As an indirect approach, we identify disintermediation using job outcomes that might imply jobs that ended prematurely or with partial payment. We collect the number of working hours for each assignment and denote it as *Hours*. For this variable, we have observations only for hourly jobs. Similarly, jobs with small payments could signal disintermediation if clients paid only a small amount on the platform and conducted most of the transaction off-platform. We use *Total_Charge* to represent the total amount paid on the platform once a job is closed.

As a direct approach to measuring disintermediation, we quantify clients' and freelancers' intentions to disintermediate by leveraging a text-analysis tool the platform developed to detect sensitive words that imply disintermediation. The list of sensitive words along with their relative weights was developed by the company based on extensive data from past transactions. Over the years, the platform refined its algorithm and dictionary based on data collected from actual disintermediation and from its other trials that were aimed at deterring disintermediation.[32] Table

[31] By design, an SS is only available after a freelancer has sufficient historical job data on the platform, which is usually after five projects or having worked with at least three clients. Freelancers without SSs appear the same to clients in the treatment and control groups and thus are not considered part of the study.

[32] For example, the platform learns about disintermediation when users seek its help with disputes or payment enforcement for transactions off the platform. In other trials, the platform issued warning messages based on

12 presents examples of sensitive words and phrases. For each message associated with a given job assignment, we sum up the numeric values of sensitive words and use the maximum value among all messages as the disintermediation score for that assignment (*Disintermediation_Score*). Compared to approaches that add up sensitive keywords in all messages or that take an average of all messages, our approach has two advantages. First, users who communicate more are likely to use more sensitive keywords; our measure is independent of the frequency of communication. Second, because users typically express their desire to disintermediate in only a few sentences and hence not all messages are useful for detecting disintermediation, our approach allows us to focus on messages that are most likely related to disintermediation. Out of the 33,561 assignments, 29,690 have historical messages, for which the platform attempts to detect sensitive words that imply users' intent to disintermediate. For job assignments whose messages have no sensitive words, *Disintermediation_Score* is 0; otherwise, it is a positive integer.[33]

Table 12: Examples of Sensitive Words/Phrases Indicating Disintermediation in Messages

off (the platform's name)	Paypal	Venmo
wire me/you/us	avoid fees	apply at/here
outside (the platform's name) / outside of (the platform's name)		
save 10% (or 5%) / save 10 (or 5) percent / save ten (or five) percent		
(my/your/our) (phone / number / phone number / cell phone)		

For each assignment, the dummy variable *Treated* equals 1 if the client is in the treatment group and 0 otherwise. To account for different levels of trust among high-quality vs. low-quality

disintermediation scores of chat messages and has improved the accuracy of disintermediation scores based on user feedback.

[33] Our results continue to hold if we let *Disintermediation_Score* = 0 for job assignments with no messages.

freelancers, we create a dummy variable, *SS_High,* which is 1 if SS $\geq$ 90%, based on the fact that the platform explicitly informed the clients in the treatment group that an SS above 90% is considered "excellent."[34]

Table 13 provides summary statistics for all variables. The unit of analysis is a job assignment. The average SS in our sample is 0.739 and the mean for *SS_High* is 0.364, meaning that 36.4% of the freelancers hired by clients have an SS of 90% or higher. In contrast, the five-star ratings of freelancers have a mean of 4.77 and a median of 5. Because the five-star ratings do not contain sufficient variation to discern freelancer quality, unsurprisingly, the SS has a low correlation with five-star ratings (0.284). The two indirect measures for disintermediation, *Hours* and *Total_Charge*, have substantial variation across job assignments and are highly skewed. The mean *Disintermediation_Score* is 7.528, with a maximum value of 34.1 and the distribution is skewed, thereby suggesting that most users in our sample have a relatively low tendency to disintermediate. Given the skewed distributions, we use the logarithms of these three measures in our regression analysis.[35] Among all assignments in our sample, 8% of the client-freelancer pairs are located in the same country as each other.

[34] As a robustness check, we replace the threshold for *SS_High* with 75%. All our findings still hold. The use of this dummy variable allows us to take the non-linear effect of SSs into account.

[35] We add 1 to *Total_Charge* and *Disintermediation_Score* before taking logarithms to avoid taking logarithms of zeros.

Table 13: Summary Statistics and Correlations

Variable	Observations	Mean	Std. dev.	Min	Max	(1)	(2)	(3)	(4)	(5)
(1) Treated	33,561	0.504	0.500	0	1	1				
(2) SS	33,561	0.739	0.258	0	1	0.041	1			
(3) SS_High	33,561	0.364	0.481	0	1	0.034	0.679	1		
(4) Hours	14,593	109.00	349.0	0.183	11365.38	-0.003	0.058	0.081	1	
(5) Total_Charge	33,561	679.507	4198.7	0	195731.5	0.005	0.049	0.062	0.683	1
(6) Disintermediation_Score	29,690	7.528	5.642	0	34.1	0.027	-0.059	-0.051	-0.021	0.003

Notes: The number of observations for the main analysis sample is 33,561, except for regressions with *Disintermediation_Score*, which has a non-missing value for 29,690 observations. *Hours* has values only for hourly jobs.

To confirm the randomness of assignment into either group, we compare the transaction data of clients in the treatment and control groups from the six-month period immediately preceding the study. The balance check, shown in Table 14, confirms that the assignment is indeed random.

Table 14: Comparison of Clients in the Treatment and Control Groups before the Study

Outcome Variable	Treatment		Control		Paired t-test
	Mean	Standard Error	Mean	Standard Error	t-stats
# of days on the platform	611.17	4.62	605.03	4.63	-0.94
# of jobs in the past 6 months	2.60	0.09	2.73	0.15	0.76
Avg. past job feedback	4.81	0.01	4.81	0.01	0.55
Avg. past job hours	12.62	0.41	13.05	0.44	0.71
Avg. past job total charge	197.00	7.77	185.55	5.31	-1.21

Notes: The unit of analysis is a client in the treatment/control group. Variables are calculated using a past assignment sample including participating clients' job outcomes in the 6 months before the study. None of the above paired t-test results is significant. We also checked the percentage of hourly jobs across the two groups and found no significant difference; the numbers are unreported due to protection of confidentiality.

Table 15 compares the three key dependent variables for the treatment and control groups after the treatment. We divide our sample by *SS_High*. The values of the three dependent variables for high-SS assignments differ significantly between the treatment and control groups. The treated job assignments have fewer total hours, lower total charges, and higher disintermediation scores, all of which suggest a greater likelihood of disintermediation. For low-SS assignments, we observe no significant differences in the three dependent variables between the two groups.

Table 15: Comparing Treatment and Control Group Observations after Treatment

Outcome Variable	Treatment		Control		Paired t-test
	Mean	Standard Error	Mean	Standard Error	t-stats
Assignments with high freelancer SSs:					
Log(Hours)	2.95	0.04	3.09	0.04	2.80***
Log(Total_Charge)	4.67	0.02	4.76	0.03	2.61***
Log(Disintermediation_Score)	1.76	0.01	1.68	0.01	-4.52***
Assignments with low freelancer SSs:					
Log(Hours)	2.66	0.03	2.67	0.03	0.09
Log(Total_Charge)	4.26	0.02	4.26	0.02	0.10
Log(Disintermediation_Score)	1.83	0.01	1.81	0.01	-1.47

Notes: The unit of analysis is a job assignment from a treatment/control group client during the study. Variables are calculated using the assignment sample for our main analysis. All the paired t-test results in the high-SS group are significant. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

II.4. Empirical Results

II.4.1. Job Fill Rates and Platform Revenue

We first analyze the impact of disclosing SSs on the job fill rate. We begin by comparing the job fill rate in the treatment group versus the control group for all job openings posted by clients. We compute the job fill rate as the ratio of the total number of filled jobs for a group to the total number of job openings for that group during the study. We also calculate the number of days before a job opening is filled. We find that the number of days taken to fill a job is 0.48% shorter for the treatment group than the control group, and the fill-rate difference is 0.51%; neither difference is statistically significant.[36] In addition, a similar number of jobs are posted for each group, with the clients in the treatment group posting only 1.3% more jobs than the clients in the control group clients; a t-test mean comparison reveals that the average number of jobs offered by each client does not differ significantly between the treatment and control groups ($p = 0.88$).

These results indicate that revealing SSs does not have a significant effect on job postings or job fill rate, which is consistent with the intuition that, when there is a sufficient number of freelancers, the job fill rate will not increase with better reputation measures.

Further, we also compare the characteristics of the jobs posted between the two groups. We find that there are no significant differences in the distributions of jobs posted across categories ($p = 0.204$ from a Chi-squared test), in the length of job descriptions ($p = 0.303$), and in clients' posted job prices ($p = 0.553$) across the two groups. These results—along with the findings that the days to fill, fill rate, and the average number of jobs by each client are not significantly different—suggest that clients' needs are exogenously determined and are not influenced by the availability of SSs. Therefore, our findings are unlikely to be driven by job-level differences between the two groups.

Next, we investigate whether revealing SSs affects the probability of hiring high- versus low-quality freelancers. For this, we obtain data on all freelancers who submitted proposals for jobs posted by clients in the study. Simple summary statistics show that the percentage of high-SS freelancers hired in the treatment group is 4.1% higher than in the control group. Then, we use a linear probability model regressing a *Hired* dummy on *SS_High*, *Treated*, and the interaction between them. The regression results shown in Table 16 indicate that freelancers with higher SSs

[36] As requested by the company, the actual values of the measures are not reported in order to protect the company's data confidentiality.

are significantly more likely to be hired than freelancers with low SSs and that revealing SSs increases the likelihood of a high-quality freelancer being hired by 0.48%, from a base of 2.8%.

Table 16: Linear Probability Model of the Treatment Effect on Freelancers' Probability of Being Hired

Model	(1)	(2)
Dependent Variable	Hired	Hired
Treated	0.0004	-0.0012***
	[0.0004]	[0.0004]
SS_High	0.0079***	0.0049***
	[0.0004]	[0.0006]
Treated x SS_High		0.0060***
		[0.0008]
Observations	895,882	895,882
R-squared	0.0004	0.0005

Notes: The unit of analysis is an application to a job posted by the treatment/control group client during the study. The mean for the dummy variable *Hired* is 0.029; the standard deviation is 0.169. Robust standard errors in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Finally, we examine the impact of the treatment on the platform's revenue. The results of the two-sample t-test show that the average revenue from each job does not significantly change for the treatment group as compared to the control group (the ratio between the treatment and control groups is 0.996). This result holds regardless of job type. We also check the percentage of successful jobs for the two groups—jobs that were completed with no abnormal status, such as "Inactivity," "No response," or "Cancelled"—and find that there is only a 0.41% difference in the percentages of successful jobs, which is not statistically significant.

These findings raise an interesting question: while revealing SSs leads to increased hiring of high-quality freelancers who often command higher prices (in the control group, the average charge for jobs involving high-SS freelancers is 2.23 times that of jobs involving low-SS freelancers), and given that job fill rates are the same, why does the platform not earn more revenue from the treatment group than from the control group? Next, we provide evidence that disintermediation, not the selection of freelancers, is the key factor that offsets the potential gain.

II.4.2. Evidence of Disintermediation

We investigate whether clients in the treatment group are more likely to disintermediate when freelancers' SSs are high, using the following regression specification at the job-assignment level:

$$Y = \beta_0 + \beta_1 Treated + \beta_2 SS\_High + \beta_3 Treated \times SS\_High + \varepsilon. \quad (4)$$

Table 17 reports the regression results. *Log(Hours)* is the dependent variable in Models (1) and (2). Model (1) shows that, on average, fewer working hours are reported for jobs in the treatment group than in the control group, and more working hours are reported for jobs with high-SS freelancers than with low-SS freelancers. Model (2) shows that displaying a freelancer's SS reduces the hours reported by high-quality freelancers by 13.3% for the treatment group relative to the control group.[37] Models (3) and (4) show similar patterns using *Log(Total_Charge)* as the dependent variable. Revealing a high-quality freelancer's SS to the client decreases the total charge by 8.2% relative to the control group.

[37] We also add category controls to each model in Table 6, and the results remain virtually unchanged. We also use the difference between each job's posted price and actual price as an alternative dependent variable. Because posted prices are similar between the two groups, unsurprisingly, we find that the treatment group's actual charge is significantly less than the posted charge for jobs with high-SS freelancers, relative to the control group.

Table 17: OLS Regressions on the Treatment Effect of High SSs on Disintermediation

Model	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Log(Hours)	Log(Hours)	Log(Total_Charge)	Log(Total_Charge)	Log(Disintermediation _Score)	Log(Disintermediation _Score)
Treated	-0.058*	-0.003	-0.034*	-0.002	0.038***	0.018
	[0.030]	[0.037]	[0.019]	[0.024]	[0.010]	[0.012]
SS_High	0.348***	0.422***	0.455***	0.499***	-0.098***	-0.128***
	[0.032]	[0.046]	[0.021]	[0.030]	[0.010]	[0.015]
Treated x SS_High		-0.143**		-0.086**		0.058***
		[0.064]		[0.041]		[0.021]
Observations	14,593	14,593	33,561	33,561	29,690	29,690
R-squared	0.009	0.009	0.015	0.015	0.003	0.004

Notes: Observations are all job assignments created during the study period. The sample in Model (1) contains only hourly jobs. The sample in Models (5) and (6) contains only jobs for which there were chat messages exchanged. *SS_High* is defined as the freelancer having a Satisfaction Score greater than or equal to 0.9 at the time of the study. *Treated*, a dummy variable for treatment at the client level, equals 1 if the client is in the treatment group. *Log(Hours)* is the logarithm of the number of hours the freelancer worked on the assignment. *Log(Total_Charge)* is the logarithm of the total amount of money charged at the end of the assignment plus 1. *Log(Disintermediation_Score)* is the logarithm of the disintermediation score computed from all messages associated with the assignment plus 1. Robust standard errors in brackets. *** p<0.01, ** p<0.05, * p<0.1.

One might be concerned that the higher probability of clients in the treatment group hiring high-quality freelancers would confound our results. However, we find no correlation between a freelancer with a high SS and the freelancer charging less or working faster (see Section 5.2). Further, if clients in the treatment group are matched with freelancers who are more productive and charge less, they should have lower incentives to disintermediate. In Models (5) and (6), we use *Log(Disintermediation_Score)* as the dependent variable. We find that clients in the treatment group are significantly more likely to disintermediate when the freelancer has a high SS.[38]

Overall, the evidence suggests that increased trust leads to more disintermediation. As a result, although providing SSs makes clients more likely to work with high-quality freelancers, the expected revenue increase is offset by disintermediation.[39]

We perform a back-of-the-envelope analysis to estimate revenue loss due to disintermediation. If the treatment were rolled out to our control group as well, the control group's percentages of high- and low-SS freelancers hired—31.4% and 68.6%, respectively—would switch to the treatment group's 35.5% and 64.5%. In other words, the proportion of high-quality freelancers hired by the control group would increase by 4.1% (35.5% - 31.4%). The average total charge for a job by a high-SS freelancer in the control group is 2.23 times that of a job by a low-SS freelancer in the control group; thus, replacing a low-SS freelancer by a high-SS freelancer on the same job in the control group would increase the job's total charge by 223% - 1 = 123%. Putting all of this together, the hypothetical rollout of SSs to all clients would create 4.1% more

[38]As the company-assigned weights for the sensitive words could affect the accuracy of the *Disintermediation_Score*, we also run a robustness check by repeating the main analysis using a dummy dependent variable to indicate whether the score is above the median score. We find that our results continue to hold.

[39]To demonstrate that SSs are indeed superior to five-star ratings, we repeat the analysis separately for freelancers with high SSs and high five-star ratings and freelancers with high SSs but low five-star ratings. We find similar results in each case, thereby suggesting that once SSs are offered, clients depend much less on five-star ratings (see Appendix Table A7).

jobs by high-SS freelancers, each generating 123% more revenue; thus, if it were not offset by disintermediation, the total revenue should increase by approximately 123% × 4.1% = 5.0% when SSs are introduced.

Note that there is a certain level of disintermediation even in the control group. Our estimated revenue loss represents only the additional revenue loss due to disintermediation beyond the baseline.

II.4.3. Heterogeneous Tendencies to Disintermediate

We examine various factors that moderate the impact of trust on disintermediation. Appendix Table A6 provides the summary statistics of all moderators.

Geographical proximity. When the client and freelancer are in the same country, they tend to have similar cultural backgrounds that allow them to build trust more easily than those from different countries. Proximity also reduces the cost of collaborating outside the platform, since they may have many other convenient channels for payment and interaction. Thus, the impact of SSs on disintermediation should be higher for jobs involving clients and freelancers from the same country. We create the dummy variable *Same_Country*, which equals 1 when clients' and freelancers' user profiles show them to be in the same country, and 0 otherwise.

Model (1) of Table 18 reports the regression results using our direct measure, the logarithm of *Disintermediation_Score*, as the dependent variable and including *Same_Country* as the moderator.[40] The coefficient for the three-way interaction suggests that clients and high-quality freelancers from the same country are indeed more affected by the treatment.

---

[40] The results obtained using the two indirect measures, *Log(Hours)* and *Log(Total_Charge)*, as the dependent variables were qualitatively similar.

Table 18: Heterogeneity in Disintermediation Tendencies

Model	(1)	(2)	(3)	(4)
Treated	0.019	0.006	0.025	0.063***
	[0.013]	[0.023]	[0.023]	[0.017]
SS_High	-0.102***	-0.126	-0.131***	-0.124***
	[0.016]	[0.028]	[0.029]	[0.021]
Treated x SS_High	0.041*	-0.026	0.028	0.021
	[0.022]	[0.039]	[0.040]	[0.030]
Same_Country	-0.037			
	[0.032]			
Treated x Same_Country	-0.013			
	[0.046]			
SS_High x Same_Country	-0.242***			
	[0.051]			
Treated x SS_High x Same_Country	0.154**			
	[0.074]			
Divisible_Med		0.079***		
		[0.019]		
Divisible_High		0.243***		
		[0.030]		
Treated x Divisible_Med		0.011		
		[0.028]		
Treated x Divisible_High		0.022		
		[0.042]		
SS_High x Divisible_Med		-0.000		
		[0.034]		
SS_High x Divisible_High		0.048		
		[0.053]		
Treated x SS_High x Divisible_Med		0.107**		
		[0.047]		
Treated x SS_High x Divisible_High		0.144**		
		[0.072]		
Long_Term			-0.114***	
			[0.026]	
Treated x Long_Term			-0.038	
			[0.040]	
SS_High x Long_Term			-0.021	
			[0.046]	
Treated x SS_High x Long_Term			0.112*	
			[0.063]	
Client_Rating_High				-0.132***
				[0.017]
Treated x Client_Rating_High				-0.089***
				[0.024]
SS_High x Client_Rating_High				-0.012
				[0.030]
Treated x SS_High x Client_Rating_High				0.071*
				[0.041]
Observations	29,690	29,690	12,118	29,690
R-squared	0.006	0.014	0.010	0.014

Notes (Table 18): The dependent variable in this table is *Log(Disintermediation_Score)*. Observations are the job assignments created during the study with non-missing disintermediation scores. Model (3) includes only jobs with information on the expected duration. Robust standard errors in brackets.

Job divisibility. The tendency to disintermediate may also vary for different job categories. Certain job categories are modular by nature and can be divided into independent parts. If a job is modular—that is, if it is more likely to be divided into parts without affecting the overall quality of the outcome (e.g., Baldwin and Clark 2000)—it is easier to perform a portion of the job on the platform, have the client check the output, and then complete the remainder off the platform. Thus, we expect the impact of treatment to be greater for more divisible jobs.

Two types of jobs are considered to be more divisible than others: (a) hourly jobs and (b) fixed-price jobs with more than one hired freelancer. We compute the percentage of such jobs in each of the platform's 13 job categories and rank the categories by that percentage from high to low; the percentages range from 33% to 88%. We create dummy variables to place all job categories into three groups based on the extent to which the jobs are divisible: the *Divisible_High* = 1 group, the *Divisible_Med* = 1 group, and the baseline group. *Divisible_High* equals 1 for the three categories with a percentage of divisible jobs in the top 10% of the divisibility distribution: Customer Service, Sales & Marketing, and Accounting & Consulting. The benchmark group includes the job categories with the least divisible jobs, with a percentage of divisible jobs in the bottom 10% of the distribution: Translation and Design & Creative. For the remaining categories, whose divisibility is between the top and the bottom 10% of the distribution, *Divisible_Med* equals 1.[41]

[41] We use category information to define job divisibility because, as we show in Section 5.7, clients may strategically choose their job types after their first jobs. As a robustness check, we use job type information from jobs posted within one month prior to the RCT and obtain the same sets of categories in each of the three groups.

Model (2) of Table 18 reports the results. As expected, the increase in disintermediation scores for clients in the treatment group who hired high-quality freelancers is greater for divisible jobs.

Expected duration. Jobs that last a long time tend to have higher costs, thereby generating the most value for the platform. They also face the highest risk of disintermediation because clients and freelancers have the greatest incentive to take the transaction off the platform. Therefore, we expect to see a greater impact of the treatment on long-term versus short-term jobs.

We use the expected job duration, which is selected by the client when posting the job and ranges from "less than one week" to "more than six months," to create a dummy variable called *Long_Term*. This information is self-reported regardless of the job type. For any job that reports an expected duration, *Long_Term* equals 1 when that duration is over six months and 0 when it is six months or less.

As Model (3) of Table 18 shows, long-term jobs indeed have a significantly higher disintermediation score for treated clients and high-quality freelancers as compared to short-term jobs.

Client rating. Having established that the disintermediation tendency varies with the business reputations of freelancers, we investigate whether it also varies with clients' reputations. We expect that, given the same SSs, freelancers tend to trust highly-rated clients more and are, therefore, more willing to take the job off-platform.

Using a client's past transactions and corresponding star-rating feedback, we compute the number of five-star jobs in each client's job history. We create the dummy variable *Client_Rating_High*, which equals 1 when that fraction is higher than the median for all clients, and 0 otherwise. Model (4) of Table 18 reports the results. We find that, when both the freelancer and the client are trustworthy, revealing more information about the freelancer's business reputation boosts their willingness to work off-platform.

We also examine client size as a possible factor that could lead to heterogeneous tendencies in disintermediation on the platform, using self-reported client-size data. As large companies are usually less concerned about cost and more about quality, they tend to place more value on the platform's role in facilitating transactions and, thus, have less incentive to disintermediate, even with trustworthy freelancers. We expect individual clients or clients who post jobs for smaller companies to be more sensitive to job cost than clients who post jobs for large companies. We do not find significant results. The lack of significance could be due to insufficient data, since the information on firm size is self-reported and is missing for 97.4% of the jobs in our sample. It could also be that, since the service fee is a fixed percentage of the transaction value and can become substantial for large jobs, large clients may find the savings from disintermediation just as attractive as small clients do.

II.5. Robustness Checks

II.5.1. Contamination Between the Treatment and Control Groups

As with any experiment conducted in a real platform, one might be concerned about violations of the stable unit treatment value assumption (Blake and Coey 2014). This concern of possible contamination is mitigated by two factors. First, the total number of clients in the trial constitutes only 3% of all clients on the platform. Second, the platform has substantially more freelancers than clients. Indeed, the job fill rates for the two groups are not different.

Nevertheless, we consider two possible sources of contamination. First, in our setting, treated clients may be better able to target high-quality freelancers, thereby leaving the control clients a pool of lower quality prospective job applicants. This may result in an exaggeration of the treatment effect when we compare the quality of freelancers hired by clients in the two groups. We conduct a t-test to compare the mean SSs of freelancers who apply to job posts in each group and find that there is no significant difference between the SSs of applicants in the treatment group and those in the control group ($p = 0.49$).

Second, it is possible that some freelancers, after the disclosure of their SSs, are approached more often by clients in the treatment group to disintermediate the platform, and may in turn suggest to clients in the control group to conduct transactions off the platform. This possibility may result in an underestimation of the treatment effect. We find that 11.1% jobs in the control group have been assigned to freelancers who have been matched to clients in treatment groups. After excluding these jobs, as expected, our results become slightly stronger (Appendix Table A8).

II.5.2. Selection of Freelancers

Prior to the availability of SSs, clients may have used price as a quality signal and selected freelancers who charge more. However, with the availability of SSs, clients may select freelancers who have a high SS but charge less because they are more efficient and can complete jobs faster. While this scenario does not explain our findings based on *Disintermediation_Score*, it is consistent with our findings that jobs done by high-SS freelancers in the treatment group take less time and cost less.

To test this alternative explanation, for each freelancer, we calculate the average number of job hours and average total charge for all assignments completed in the six months prior to the study timeframe; we use *Past_Hours* and *Past_Total_Charge* to denote these variables. We then repeat our analysis in Models (1) through (4) of Table 17 with the logarithms of *Past_Hours* and

Past_Total_Charge as dependent variables. If high-SS freelancers were hired by clients in the treatment group because they could work more efficiently and, therefore, charged less than freelancers in the control group, this difference should be evident in their work completed prior to the study.

Appendix Table A9 reports the results. We find that high-SS freelancers hired by clients in the treatment group during the study do not appear to work faster or charge lesser than high-SS freelancers in the control group prior to the study. The results also suggest that clients in the treatment group are not selecting freelancers with greater tendencies to disintermediate.

As another robustness check, we compare the actual job duration—measured by the number of days between the start date of a job to the date on which the client marks the job as completed—to the expected job duration when the client posts the job. A client is asked to indicate estimated job duration when posting a job by selecting one of the five following options: "More than six months," "Three to six months," "One to three months," "Less than one month," and "Less than one week." First, we compare the distribution between the treatment and control groups and do not observe significant differences in expected job durations ($p = 0.239$), thereby suggesting that these jobs are comparable. If the fewer hours for jobs in the treatment group are indeed caused by clients' selection of more efficient freelancers, we would expect the actual duration of these jobs to be shorter. However, if they are caused by disintermediation instead, the actual job duration may not be shorter. By not marking the jobs as completed before their actual completion, the clients could leverage instrumental trust (Bapna et al. 2017) by reserving the option to report to the platform or leave negative ratings if the freelancers do not complete the jobs satisfactorily.[42] We

[42] The platform allows clients and freelancers to leave each other reviews within 14 days after their jobs are marked as completed.

compare (the logarithm of) the actual job duration between the two groups, and find that jobs in the treatment group take 0.75% longer than those in the control group ($p = 0.260$).

These results boost our confidence that the differences in the hours and total charges are caused by greater disintermediation.

II.5.3. Robustness of the Disintermediation Score Measure

Because the disintermediation score is constructed from a dictionary of company-developed keywords, one might be concerned about whether the keywords accurately capture user intention to disintermediate. For example, when there is a great match between a client's need and a freelancer's skill, but the freelancer has a very low SS. Before hiring such a freelancer, the client may devote more efforts into interviewing the freelancer to mitigate the risk, which may involve asking for the freelancer's contact information. However, such behavior does not indicate intention to disintermediate. Although these cases bias against finding an effect,[43] to address this concern, we create a more robust disintermediation score based only on keywords that explicitly indicate an intention to disintermediate, such as "avoid fees," "save 10%," and "wire me/you/us," and excluding ambiguous keywords such as "your phone number" and "Skype."

Appendix Table A10 reports the results after repeating our analysis with this new disintermediation score (*Disintermediation_Score_Robust*). The results obtained are qualitatively the same. The coefficient of the interaction term is smaller because the new dictionary includes fewer keywords.

[43] This also helps explain why the coefficient for the "Treated" variable is positive in Model 6 of Table 6.

II.5.4. Is it All About Speeding up the Inevitable Outcome?

Tryouts between clients and freelancers may also help build trust (e.g., Gulati 1995). After tryouts, clients and freelancers who had positive experiences with one another may decide to disintermediate. If disintermediation is inevitable because of tryouts, introducing a better reputation score may only help speed up this outcome by reducing the duration of the tryout period.

However, tryouts have their limitations and are not perfect substitutes for a better reputation system. First, for a durational job, the experience can vary within the same job—it is not uncommon to have a good start and a terrible ending. In other words, the initial experience in a job does not necessarily indicate its final success.

Second, experiences with the same freelancer can vary across jobs. For example, it is possible for freelancers to build their reputation first and later exploit customers who trust them the most (e.g., Liu 2011). Further, clients typically use online freelance platforms to fulfill ad-hoc needs. On the platform that we studied, although the platform has been in the business for more than ten years, the median number of successfully filled jobs per client is two.[44] Thus, for most clients, their limited number of past interactions would not sufficiently help reduce uncertainty.[45] These clients would value an accurate reputation system—which includes freelancers' past performance for many clients—because these ratings capture the overall satisfaction for all jobs completed, and the rating variation represents truthful uncertainty.

Consistent with these arguments, Kim et al. (2012) find that trust exerts a stronger effect than perceived price on purchase intentions for *both* new and repeat customers of an online store.

---

[44] A similar statistic is also reported by Barach et al. (2018) in their study of a similar freelance platform.

[45] Because our sample collection requires a client to post a job during the 26-day trial, the RCT selects clients that post jobs more frequently.

In a setting like ours, where tasks are heterogeneous across jobs and over time within a job, we expect that the efficiency of tryouts in building trust is even lower.

In fact, one might expect that SSs have a greater impact on clients who have had some past interactions with the same freelancers. A good SS and clients' personal positive experiences can work in combination to build sufficient trust for clients to be willing to disintermediate. Furthermore, after having some interactions with freelancers, clients may feel more comfortable suggesting taking transactions off the platform.[46]

In 21.1% of the job assignments in our data set, the clients have already had past interactions with the same freelancers. These clients most likely had positive experiences with the freelancers and thus decided to work with them again. If we restrict the analysis to these repeated interactions, we find that SSs continue to have significant effects (Appendix Table A11). The magnitude of the effects is actually greater for repeated hires than for first-time hires, thereby suggesting that even with tryouts, certain clients still choose not to disintermediate and that the introduction of SSs, in addition to their own experiences, motivates them to disintermediate.[47]

II.5.5. Client Satisfaction with High-SS Freelancers

Another potential explanation for the lower fees and hours is that high SSs may have raised treated clients' expectations of the freelancers. Treated clients may therefore end jobs prematurely and pay less as a result of dissatisfaction rather than disintermediation.

[46] As in most marketplaces, disintermediation is a violation of the terms of services on this platform, and the platform encourages users to report such behavior.

[47] Note that the results from Model (3) in Panels A and B of Appendix Table A11 may appear to suggest that job assignments with high-SS freelancers are less likely to be disintermediated in repeated relationships than during first-time hires. We obtain these because disintermediation is harder to detect using our disintermediation score for repeated transactions. The client and the freelancer could have already exchanged contact information when they worked for the first time. They could disintermediate without exchanging the information again during repeated transactions.

To test this alternative explanation, we collect data on clients' feedback on freelancers after each job assignment in our sample. The number of jobs for which the client did not leave any feedback is similar for both the treatment and control groups (30.1% and 29.8%, respectively). We replicate the analysis in Table 17, replacing the dependent variable with the client rating for each job assignment. Appendix Table A12 shows that neither the coefficient for being in the treatment group in Model (1) nor the coefficient of the interaction term in Model (2) are statistically significant. Thus, our findings are not driven by reduced client satisfaction with the work of high-SS freelancers.

II.5.6. Who Initiated the Disintermediation?

As SSs were only revealed to the clients in the treatment group, if revealing SSs indeed resulted in more disintermediation, we expect that clients in the treatment group are more likely to initiate disintermediation. Among the job assignments in which disintermediation is detected, a t-test comparison shows that clients in the treatment group indeed initiate disintermediation 3% more frequently as compared to clients in the control group ($p = 0.042$).

II.5.7. Clients' Strategic Choice of Job Type

After observing SSs, clients interested in disintermediation might strategically select hourly jobs, as these are more conducive to disintermediation. We test this using data on clients' job openings and running several t-tests to compare the number of hourly versus fixed-price jobs that clients posted over time.

Overall, the distribution of job types is different between the treatment and control groups ($p = 0.029$), with clients in the treatment group posting a slightly larger proportion of hourly jobs than clients in the control group (1.2% higher). We split the job-opening sample into subsamples of clients' first job posts and subsequent job posts and then repeat the comparison for each subsample. We find no difference in the distributions of job types between the treatment and control groups in the subsample of clients' first jobs ($p = 0.986$). However, there is a significant difference for the subsample of subsequent job posts: the percentage of hourly jobs posted is 3.5% higher in the treatment group ($p = 0.0004$).

This result suggests that while clients in the treatment group do not appear to take SSs into account when they post their first job (as most become aware of SSs only when they review proposals for their first job posts), they are more likely to post hourly jobs subsequently. This result supports the explanation that clients in the treatment group are, on average, more inclined to disintermediate.

II.6. Discussion and Conclusion

We provide empirical evidence on disintermediation and show that disintermediation can sometimes render less effective a platform's strategy to improve its profitability through enhancing trust. An important challenge in studying disintermediation is the observation of disintermediated transactions. Our study shows that digitization can help overcome this challenge by capturing detailed data on user communication and activities.

We examine only one type of disintermediation. In reality, there are other means by which that users can disintermediate an online platform. For example, a user can use a platform to find an ideal match and then directly contact the other party without ever initiating a transaction on the platform. In this study, the fill rate does not significantly decrease for clients in the treatment group, but this type of disintermediation could occur more frequently in other settings. A user may also complete one transaction on a platform and then take all future transactions with that party off the platform.

In addition, our study examines only the short-term effect of building trust. As more users on the platform realize the benefits of disintermediation, there could be an increase in the negative effects of enhanced trust on platform revenue. Our result that clients change their job format after posting their first job suggests that they do learn from their experiences. Since not all clients hired freelancers twice during our experiment, we would expect to have more disintermediation as a result of this learning if the experiment had run for a longer period of time. Overall, the strategic behavior may increase the negative effects of trust building in the long term. At the same time, a better trust-building mechanism may attract more users to the platform, or may incentivize clients to post more jobs or jobs of higher values, thereby resulting in greater revenue. Thus, the long-term effect is ambiguous.

Finally, while we observe repeated interactions in our data, our experiment does not allow us to estimate the causal relationship between repeated interactions and disintermediation. With repeat interactions, one needs to worry about reverse causality: whether repeated interactions lead to disintermediation, or whether the intention to disintermediate lead to repeated interactions. Future research could conduct more experiments to compare the effectiveness of a better reputation score and repeated interactions or other trust-building mechanisms in driving disintermediation.

Notwithstanding these limitations, it is important to note that the main objective of our research is not to conduct a net benefit analysis to determine whether it is optimal for platforms to implement SSs. Rather, because trust building is important for platform growth, it is of vital importance for a platform to build as much trust as possible. At the same time, our research highlights a negative effect of trust building and suggests that as a platform builds more trust to facilitate transactions in its marketplace, it needs to adopt appropriate strategies to counter increased disintermediation.

Platforms could use a variety of strategies to reduce disintermediation as they enhance trust. Airbnb, for example, enhances trust and safety through host ID verification and background checks. At the same time, Airbnb reduces disintermediation by withholding host data, such as listing address or phone number, until the payment is made. Thumbtack, a marketplace that connects consumers with local service providers such as house cleaners, captures value pre-transaction: when customers post job requests on Thumbtack, service providers can send quotes to the customers; service providers pay fees to Thumbtack only if customers respond. Disintermediation affects Thumbtack less strongly because its model captures value before two parties agree to work together.

Other platforms recognize that the motivation to disintermediate comes from the service fees they charge and adopt different value-capture strategies to prevent disintermediation while still enhancing trust. For example, Chinese outsourcing platform ZBJ, which launched in 2006 with a 20% commission model, began pursuing other revenue sources after calculating that it could lose as much as 90% of its business to disintermediation. In 2014, ZBJ leveraged big data analytics to find that new business owners often used ZBJ to outsource logo design. However, after logo design, many of these clients would also need business and trademark registration. Thus, ZBJ began offering this service and has now become the largest provider of trademark registration in China. Replicating this experience, ZBJ began providing several other services to its platform participants. With these revenue streams, the company decided to significantly reduce its commission to as low as 2% and shifted its resources from fighting disintermediation to growing its user base and building trust (for example, by encouraging clients and freelancers to communicate) (for more details, see Zhu et al. 2018). Because of these changes, the company obtained a valuation over $1.5 billion in 2018. Future research could examine the effectiveness of various strategies that platforms use to mitigate disintermediation.

CHPATER III:

The Great Firewall of China and Marketplace Disintermediation

III.1. Introduction

Online intermediaries provide matching and facilitate transactions between sellers and buyers (e.g., Parker and Van Alstyne 2005; Rochet and Tirole 2006; among many others). However, the rising trend of disintermediation—the action of sellers and buyers circumventing a platform to transact directly—prevents the platform from capturing value in this process (e.g., Edelman and Hu 2016; Gu and Zhu 2018). As a result, online businesses, such as Homejoy, have failed to sustain their revenue in the long run.[48] Understanding the factors that lead to disintermediation is critical to platform value creation and value capture.

An important factor coinciding with the surge of disintermediation is the rapid development of online technology. For instance, as the Internet and digital payment technology allow people to freely transact online, authors can now directly sell to readers on Amazon instead of via traditional book publishers. Also, as web and data management services advance quickly, hotels and airlines incentivize customers to book directly on their official websites, threatening the business model of online travel agencies (OTAs).

Despite the surge of disintermediation and online technologies, it remains unclear how technology affects disintermediation and the implication for online platforms. On the one hand, improved technology enables a platform to provide enhanced services, such as payment security protection and timely job monitoring procedures, which increase the platform's value creation and incentivize users to complete more transactions on the platform. On the other hand, the ease of technology use may simplify user communication and build direct user trust, weakening the platform's role in facilitating transactions and leading to revenue "leakage" as users take

[48] About 90% of transactions involve disintermediation after customers are matched with service providers, according to a survey on one of the world's largest freelancing marketplaces (Zhu et al. 2018).

transactions offline. Given these opposite influences, there is yet to be theoretical or empirical evidence of the overall effect of technology on disintermediation.

This research examines the impact of a particular type of technology, online communication technology, on user disintermediation tendencies in two-sided marketplaces. Communication technology enables users to contact each other directly without a platform; thus, it is necessary for initiating, as well as preventing disintermediation. To be specific, the study focuses on general online communication technology, including third-party instant messaging software, such as Skype or WhatsApp, which is widely used and acts as a complementary channel to the platform's own communication tool. I investigate whether changes in such communication technology increase or reduce user disintermediation on the focal platform, as well as their impact on transaction outcomes.

The research design leverages an exogenous shock to communication technology, the sudden Skype block in mainland China, as a natural experiment to observe its impact on user disintermediation tendencies in the setting of a major United States (US) online outsourcing platform. The impact of blocking Skype is important to this major outsourcing platform, since this platform has its own messaging system but its users also use third-party communication tools to assist with their transactions, the mostly popular of which is Skype, one of the world's leading online messaging software with more than 300 million monthly active users in 2019.[49] A survey conducted by the platform shows that Skype is also the most-used software when users try to contact each other directly. In addition, the platform has a large crowd of users from mainland China, who are affected by the Skype block in 2017. Given that, over the past decade, other major

[49] Source: https://www.statista.com/statistics/258749/most-popular-global-mobile-messenger-apps/, accessed July 2019.

third-party communication software, such as Gmail, WhatsApp, and Facebook Messenger, have all been blocked in mainland China, [50] the Skype block cuts off the major remaining communication software that allowed mainland Chinese Internet users to disintermediate with users in the US and other western countries.

The study employs matching and difference-in-differences methods to estimate the effect of the shock. First, a control group is constructed among the mainland China freelancers on the platform who are affected by the Skype block in late 2017. Given China's "One Country, Two Systems" policy in Hong Kong and Macau, the "1992 Consensus" with Taiwan, and various historical connections with Singapore and Malaysia, I construct a control group of freelancers from these five locations with similar cultural backgrounds and user characteristics to the affected group. Next, using all the job transactions completed by these two groups of freelancers from January 2017 to June 2018, the difference-in-differences analysis estimates the average treatment effect of the Skype block on the affected freelancers' job hours, job total charges, and disintermediation tendencies.

While the measure for users' disintermediation tendencies is key to the analysis, disintermediated transactions are hard to observe in practice, which also makes the topic understudied in the literature. Therefore, this study takes advantage of a text analysis algorithm to analyze users' historical chat messages on the platform. Specifically, the algorithm detects sensitive keywords in the messages that indicate users' intentions to disintermediate, assigns weights to these keywords, and quantifies the users' disintermediation tendencies expressed during each job as a numeric score. This text analysis algorithm has been tested and improved over time in the platform's business practice, thus providing a direct proxy for disintermediation.

[50] Source: https://en.wikipedia.org/wiki/Websites_blocked_in_mainland_China, accessed July 2019.

The study finds that restricting alternative communication technologies outside the focal platform leads to a decrease of disintermediation on the platform. After Skype was blocked by the Chinese government, the total charges per job of mainland China freelancers increases by an estimate of 27.5%, and the total number of hours per job increases by 33.3%. Meanwhile, the freelancers' intention to disintermediate decreases by 20.5%. This evidence suggests that freelancers in mainland China are less likely to take transactions off the platform after the block; the magnitude of this main effect is larger for communication-intensive jobs and for jobs posted by personal clients instead of business clients.

III.1.1. Relationship to Prior Work on Technology and Disintermediation

The above findings mainly contribute to the growing body of literature on disintermediation, as there is not yet any work focused on communication technology as a factor that leads to disintermediation. Existing work has examined disintermediation in various contexts, such as how disintermediation disrupts the market for ideas (Peukert and Reimers 2018) and how disintermediation affects product quality over time in the music industry (Waldfogel 2012). The closest prior work, from which the measurement in this study is adapted, is by Gu and Zhu (2018). However, Gu and Zhu (2018) aims to understand how user trust, elaborated in a randomized control trial, influences disintermediation in online platforms, which is distinct from this paper's focus.

This study is also of interest to scholars studying information and communication technology (ICT). Few studies have explored disintermediation as a side effect of the rapid technological development in recent years. Most of the research in this community investigates the benefit of ICT development on various business performance measures (e.g., Hinds and Kiesler 1995; Xu et al. 2014) or its impact on organizational forms (e.g., Bloom et al. 2014). There is also literature comparing ICT's effects in various country and cultural settings (e.g., Tan et al. 2014; Venkatesh et al. 2016). The threat of disintermediation, on which this study focuses, adds another lens to discussing the impacts of communication technology.

Finally, this study is relevant to the large stream of research on value creation and value capture, as theoretical and empirical evidence is lacking on how disintermediation hinders value creation from technology. Prior discussion in the field involves value creation in the context of technology innovation ecosystems (e.g., Iansiti and Levien 2004; Kapoor and Lee 2013), competitive strategies in platform-based markets (e.g., Cennamo and Santalo 2013; Shankar and Bayus 2003), and value capturing for high-technology new entrants and incumbents in network markets (e.g., Zhu and Iansiti 2012; Fuentelsaz, Garrido, and Maicas 2014) and among technologically interdependent complementors (e.g., Adner and Kapoor 2010; Boudreau and Jeppesen 2014; Pierce 2009). There is yet to be any theoretical or empirical research on how disintermediation acts as a mechanism that affects how technology creates value.

The rest of the paper proceeds as follows. Section 2 provides background on the empirical setting and The Great Firewall of China. Section 3 describes the research design and data. Section 4 presents the empirical results. Section 5 investigates a potential mechanism and conducts robustness checks. Section 6 concludes the paper.

III.2. Background

The empirical setting of this study is a major online outsourcing marketplace in the US. It is an apt setting for the study as the platform is vulnerable to disintermediation and is also sensitive to technological change. On this platform, freelancers located in mainland China were affected by the blocking of Skype in October 2017 due to Internet censorship from the GFW. The block had no prior government announcements or news media coverage. After the block, mainland China Skype users could no longer download or update the software, refill credit, group video chat, or make landline calls via Skype.

III.2.1. Empirical Setting

Observing disintermediation is key to this analysis. As noted above, the empirical setting of this study is a major online outsourcing platform in the US. Transactions on this platform are repetitive, global, and, most importantly, communication-intensive, which makes it vulnerable to disintermediation and sensitive to technological change.

To be specific, this platform has several advantages as the research setting. First, since the platform charges freelancers a service fee of between 5% and 20% of the amount they bill a client, it provides an incentive for both sides to disintermediate to avoid the fee at various stages of the transaction. A client posts a *job opening* and receives freelancer applications. When the client hires a freelancer, an active *job assignment* is initiated until the job is complete and both sides agree to close the assignment. Jobs can be either *hourly* or *fixed price*. The amount charged for a fixed price job is negotiated and set when the job assignment begins. For hourly jobs, only an hourly rate is pre-negotiated at the time of hiring, and the total amount charged equals the number of working hours times the hourly rate. In such a setting, disintermediation can happen in two ways: before a transaction starts or after a transaction starts but before it is completed. This study focuses on the latter case, which is reflected in the job outcomes.

Second, although the platform has its own online messaging system, since the jobs on the platform are communication-intensive, using third-party software may help convey specific work needs. There are a wide range of job categories on the platform, such as writing; sales & marketing; accounting & consulting; web, mobile & software development; design & creative; and administrative support. In many of these categories, jobs are not standardized, such as designing a logo. To communicate about job requirements and progress, the client can start a conversation with the freelancer by creating a one-on-one or group *room*, which stores all future messages between the room participants. While this room offers some functions, including video and audio chatting, screenshot sharing, and file attachments, it does not provide other functions, such as computer-to-phone calls, SMS messages, or group meetings. Therefore, for non-standardized jobs, users may need other software, such as Skype, for urgent requirements or to remotely monitor work quality.

Also, a large portion of jobs are long-term and repeat transactions, which allow clients and freelancers to try each other out and then disintermediate for the rest of the job. For long-term jobs, the client and the freelancer can agree upon *milestones*, at which the client pays a corresponding part of the bill each time the freelancer completes a portion of the job until the job assignment is fully complete. Such long-term jobs usually last from a week to more than six months, with regular milestones and billing cycles, which makes it easier to disintermediate in the middle of the transaction.

Finally, freelancers on the platform are located worldwide, and, hence, rely heavily on online communication software for contact. Over 1/3 of all jobs on the platform are completed by freelancers outside the US, from regions, such as Europe and Asia. Also, a high portion of domestic clients tend to hire international freelancers. The global setting of this platform makes online chatting the major and almost only channel to communicate and monitor job quality.

In sum, the platform faces a high risk of disintermediation and is an apt setting for this study. For the above reasons, several studies have used such platforms as the empirical context to study online labor contracting (e.g., Stanton and Thomas 2016) or disintermediation (e.g., Gu and Zhu 2019).

III.2.2. The Great Firewall of China and the Sudden Skype Block

The mainland China government's Internet censorship system, known as The Great Firewall (GFW), is the world's largest online censorship system. First started in 1998, and with an affected population of 800 million Internet users today,[51] the goal of the GFW project is to legislate domestic Internet usage and regulate the Chinese Internet economy through online traffic monitoring, limited access to foreign websites, the selected blockade of foreign Internet tools and software, and so on.[52]

The GFW has restricted access to various major communication software in mainland China, including Google's Gmail and Gtalk, Facebook Messenger, Twitter, WhatsApp, and major online media sources, such as The New York Times and YouTube. The most recent target of the GFW is Skype. Skype is an online instant messaging software offering free video and audio chatting functions all over the world.[53] In 2017, Skype was listed as the world's most popular instant messaging software. In China, Skype has been widely used by people to connect to US and European online marketplaces. After other instant messaging software was blocked by the GFW, Skype became the major online messaging channel between mainland China and US Internet users. Between October 19 and October 22, 2017, the Chinese government took Skype down from all

---

[51] Source: https://www.bloomberg.com/quicktake/great-firewall-of-china, accessed January 2019.
[52] Source: https://en.wikipedia.org/wiki/Great_Firewall, accessed January 2019.
[53] Source: https://www.skype.com/en/, accessed January 2019.

major online app stores in mainland China due to "political concerns."[54] There were no prior government announcements or media coverage. The block immediately affected various Internet users in global business and online trading industries. While users could still turn to Internet proxy tools, such as the virtual private network (VPN), to circumvent the GFW, on February 1st, 2018, the Chinese government officially banned the usage of all VPNs nationwide, shutting down the last method to browse outside the GFW.

The block has had several long-lasting effects on mainland China Skype users. First, no new downloads or installations of Skype were possible for mainland China users. For existing installed versions, money refill was no longer allowed; therefore, paid services, such as landline calls and group video chatting, were unusable once the user ran out of points in his or her Skype account. Moreover, future version updates were no longer available, so current users would be forced to stop using the software, as Skype does not allow logins using older versions once it issues a major software update. With such impacts, Skype usage is likely to die out completely in China over time.

III.3. Research Design and Data

The research design uses the Skype block in mainland China in 2017 as an exogenous shock to communication technology and employs matching and difference-in-differences approaches to estimate its impact on disintermediation. The matching process is based on mainland China's government policies with Hong Kong, Macau, and Taiwan and its historical connections with Singapore and Malaysia. The matched group of freelancers show parallel pre-trends to the affected

[54] Source: https://www.nytimes.com/2017/11/21/business/skype-app-china.html, accessed January 2019.

freelancers in terms of their disintermediation tendencies and job outcomes before the shock. The main analysis uses transaction-level data for the matched pairs, including disintermediation scores and job outcomes for all job assignments completed by the affected and matched freelancers from January 2017 to June 2018.

III.3.1. Constructing a Matched Control Group

To prepare for the main difference-in-differences analysis to estimate the effect of the exogenous shock, the matching process aims to produce a control group that balanced with the treated group, but was unaffected by the block. The goal of the matching process is to provide a matched group that is comparable to the affected group before the shock to allow the estimation of the counterfactual scenario after the shock using difference-in-differences analysis.

To obtain a comparable matched group, I first take advantage of mainland China's government policies with Hong Kong, Macau, and Taiwan, since the Skype block only applies to mainland China but not to freelancers in Hong Kong, Macau, and Taiwan. Under the People's Republic of China's "One Country, Two Systems" policy in Hong Kong and Macau,[55] people living in those regions communicate frequently with mainland China, including for work, study, transport, and marriage. Similarly, given the "1992 Consensus" with Taiwan, mainland China has strong cultural and economic connections with Taiwan.

In addition, due to historical reasons, two other Asian countries, Singapore and Malaysia, have significant Chinese-speaking populations, as well as the largest overseas Chinese populations in

[55] Source: https://en.wikipedia.org/wiki/One_country,_two_systems, accessed September 2018.

the world.[56] The deep Chinese cultural and economic influences in these two countries, as well as in Hong Kong, Macau, and Taiwan, make them a highly comparable control group to mainland China freelancers.

The matching method employed is the propensity score matching (PSM), which ensures that the two groups are comparable across various exogenous covariates. The level of matching is each freelancer in the affected or control regions. Each mainland China freelancer is paired with another freelancer in the control regions with no replacement based on their likelihood of being treated. The exogenous covariates include the freelancer's days on the platform since registration, the average amount of posted charge of each freelancer's fixed price jobs when jobs are posted, the number of jobs completed, the freelancer's English language skill score based on tests on the platform, the job categories, and the job quality of the freelancer. The matching results and balance tables are described below in Section 3.3.

III.3.2. Data and Measures

This study's data set contains all job assignments completed by mainland China freelancers and their matched freelancers on the outsourcing platform from January 2017 to June 2018, providing job-level outcome data to allow for observing disintermediation tendencies per transaction. To be specific, for matching, 883 treated freelancers and 2,237 control freelancers completed at least one job assignment in 2017 before the Skype block; the matching constructs a balanced sample of 1,754 freelancers, who are equally split between the two groups. They completed 15,281 job assignments from January 2017 to June 2018, among which 7,595 job

[56] The Chinese populations in Singapore and Malaysia are 75% and 23% of the total populations, respectively. See https://en.wikipedia.org/wiki/Overseas_Chinese for a list of overseas Chinese population percentages in all countries, accessed September 2018.

assignments are completed outside the shock period between October 2017 and January 2018.[57] The clients for these jobs are limited to those located in the US to control for geographical difference on the demand side. These 7,595 job assignments form the final sample for the main analysis.

One of the key difficulties in this study is measuring the amount of disintermediation. Resembling Gu and Zhu (2019), this study employs a direct measure of disintermediation, called *Disintermediation_Score*, using text analysis based on the platform's keyword detection algorithm of chat messages sent between the clients and the freelancers.[58] To be specific, the platform constructs a list of sensitive words together with their weights based on an algorithm trained from a large set of past transaction data. For each message associated with a job assignment, the algorithm sums up the numeric values of sensitive words and uses the maximum value among all messages as the disintermediation score for the assignment. Compared to approaches that add up sensitive words in all messages or that take an average of all messages, this approach is independent of the frequency of communication; also, as users typically express their desire to disintermediate in only a few sentences, not all messages are useful for detecting disintermediation; hence, this approach focuses on the messages that are most likely related to disintermediation. The scores are learned and refined over time, and detected instances are confirmed using various other methods for detecting disintermediation during interviews or payments; thus, this measure is practically reliable. In total, 1,706 out of the 7,595 job assignments have a

[57] After Skype was blocked in October 2017, some mainland China users could still circumvent the GFW to download or update Skype using VPN or proxies. On February 1st, 2018, VPN was also banned by law. Therefore, in the main analysis, the four months from October 2017 to January 2018 are considered as the period when the shock is "being implemented," and the months since February 2018 are considered as the period "after the shock." I also did a robustness check by including October 2017-January 2018 in the "after the shock" period in the main analysis; all findings are qualitatively unchanged.

[58] The platform used a text analysis tool on aggregated data to detect conduct in violation of the Terms of Service. The actual content of chat messages was not reviewed, and all user information was kept confidential.

Disintermediation_Score that is a positive integer; the others have a *Disintermediation_Score* of 0.

Beyond the direct measure, two indirect measures on the job outcomes are also used to identify disintermediation, the number of job hours and the job's total job charge, both of which aim to measure jobs that are partially disintermediated. These two measures and *Disintermediation_Score* together quantify the tendency of user disintermediation. Specifically, as an indirect approach, jobs that are disintermediated tend to have shorter working hours, as the client only tries out part of the job with the freelancer on the platform and takes the rest outside the platform. This is captured by *Hours*, the number of working hours recorded by the freelancer time tracking software for each assignment. *Hours* is only available for hourly jobs and is missing for fixed-price jobs. Similarly, for disintermediated jobs, *Total_Charge*, the total amount paid on the platform when a job assignment is closed, tends to be small, indicating that the rest is being paid outside the platform to avoid commission fees. *Total_Charge* has a non-negative value for all job assignments.

For each assignment, the dummy variable *Treated* equals 1 if the assignment is completed by a freelancer in mainland China and equals 0 otherwise. Also, to mark the time periods before or after the Skype block, a dummy variable *Post* is assigned as 0, if the job assignment is completed between January 1st and September 30th, 2017, and as 1, if the job assignment is completed after February 1st, 2018. The time period between October 1st, 2017 and January 31st, 2018 is considered to be when the block happened; therefore, *Post* has no value in that interval, since no job assignments during that time period is included in the sample.

In addition, two moderating variables are constructed to examine the heterogeneous tendencies of disintermediation based on the communication-intensiveness of the job and the client type. First, jobs that require more communication to convey the client's specific needs rely more on video or audio chatting during the working process to monitor job quality. Therefore, such jobs are likely to be more affected by the Skype block than jobs requiring less communication. The level of communication required can be quantified by the average number of milestones per job in each job category; in other words, by whether the client constantly monitors the job quality throughout the transaction or just waits for final delivery once the job starts. The average number of milestones ranges from 5.9 to 10.2 for the 13 job categories. This allows the jobs to be split into three groups—*Communication_High*, *Communication _Med*, and the baseline group. *Communication _High* equals 1 for four categories with an average number of milestones in the top 25% of the milestone distribution: Writing, Sales & Marketing, Accounting & Consulting, and Engineering & Architecture. The baseline group includes the categories with the least number of milestones per job in the bottom 25% of the milestone distribution: Translation, Web, Mobile & Software Dev, Data Science & Analytics, and Legal. For the rest of the categories, whose number of milestones is between the top and the bottom 25% of the distribution, *Communication_Med* equals 1.

Also, jobs from business clients generally rely less on Skype for disintermediation relative to personal clients, since company employers are more likely to have their own communication tools or websites for the freelancers to contact directly. Hence, blocking Skype may affect personal clients' jobs more than those from business clients. This heterogeneity is captured by a dummy variable called *Client_Personal*, which equals 1 if the client self-report as a personal client on the user profile and equals 0 if the client self-reported as a business client.

III.3.3. Balance Check and Summary Statistics

A. Parallel Pre-Trends Assumption: To establish the validity of the casual inference in the matching and the difference-in-differences specification, it is important to investigate the parallel pre-trend assumptions of the dependent variables between the treated and matched group freelancers. Figures 12 to 14 compare the trends of the three dependent variables' changes over time for the two groups before and after the Skype block. In Figure 1, for instance, each point represents the monthly total charge per job of each freelancer in a given month, averaging among all freelancers in the treated or matched groups. If a freelancer does not complete any assignment that month, the person's total charge for that month is 0.

Figure 12: Difference between the Treated and Control Freelancer's Disintermediation Score over Time

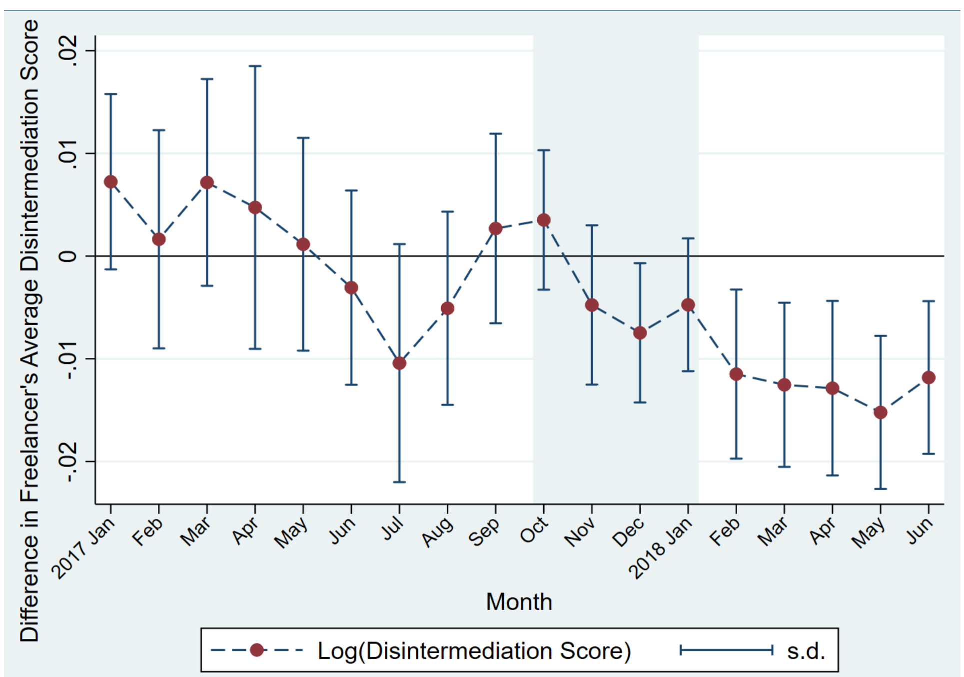

Figure 13: Difference between the Treated and Control Freelancer's Average Job Total Charge over Time

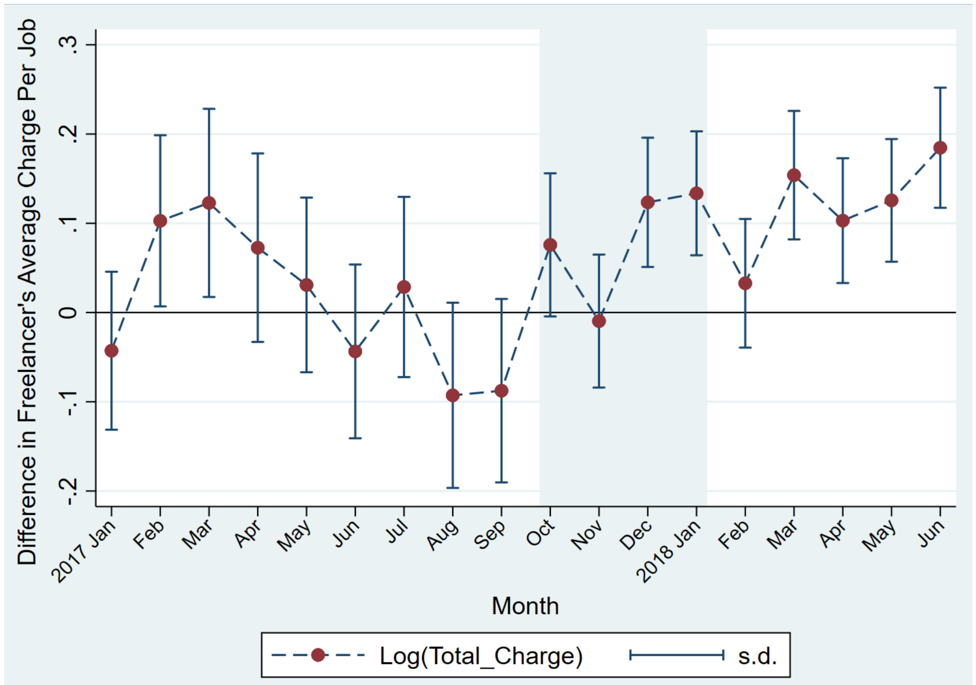

Figure 14: Difference between the Treated and Control Freelancer's Average Job Hours over Time

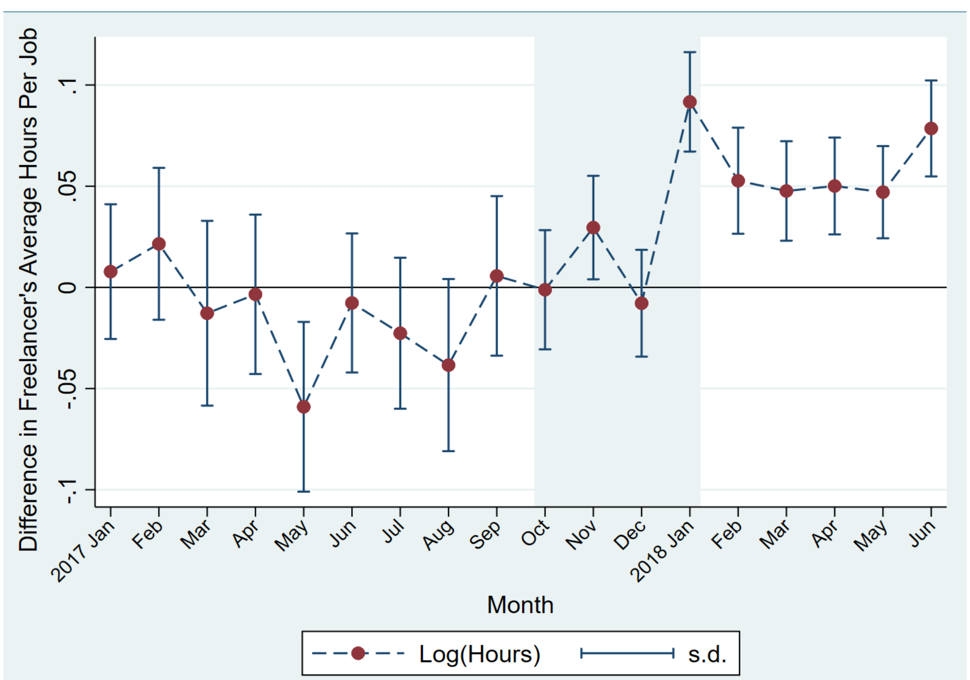

As evident from Figures 12 to 14, the pre-trends for all three variables are similar between the two groups, which verifies that the parallel pre-trends assumption is satisfied by the matching. Meanwhile, we can observe an increasing divergence between the two groups on each variable after the Skype block, visualizing the potential impact of the shock on job outcomes and level of disintermediation.

B. Balance Table: It is also necessary to make sure that the two groups of freelancers are balanced on various characteristics and dimensions before the shock. Table 19 presents paired t-tests results at the freelancer level to compare various aspects of the treated and the matched group freelancers before the Skype block, such as the freelancers' number of days on the platform, the average job opening amount in the past year, and the freelancers' language skills. We can see that the blocked and matched group freelancers are balanced on all dimensions.

Table 19: Balance Table of the Treated and Matched Freelancers' Characteristics and Job Outcomes, Before the Shock

Variables	Treated		Matched		Paired t-test
	Mean	Standard error	Mean	Standard error	t-stats
Profile Complete Percentage	86.98	0.78	88.61	0.78	1.48
Profile Available Hours	39.64	0.09	39.61	0.10	-0.39
Avg. Job Success Score	0.85	0.01	0.84	0.01	-0.54
English Level	4.92	0.01	4.94	0.01	1.40
Avg. Job Hours	3.84	0.58	4.11	0.54	0.34
Avg. Job Total Charge	122.57	15.24	131.79	17.42	0.40
Avg. Disintermediation Score	0.03	0.01	0.03	0.01	-0.29

Notes: The unit of analysis is each freelancer in the treated and control group. Variables are calculated based on the freelancers' job assignments in 2017 before the Skype block in October. The number of observations for each group is 862. None of the above paired t-test results is significant.

C. Summary Statistics: The main analysis of this study is conducted at the job assignment level. Table 20 shows the summary statistics of all variables. The number of observations in the main analysis sample is 7,595, of which 2,125 are hourly jobs. Only 1,706 jobs have non-zero disintermediation scores. Since *Hours*, *Total_Charge,* and *Disintermediation_Score* are highly skewed, the main analysis uses the natural logarithms of each of the three measures as the dependent variables. Furthermore, 48.5% of the job assignments in the main sample are from the treated group freelancers, while 51.5% are from the matched group freelancers.

Table 20: Summary Statistics

Variables	Observations	Mean	Std. dev.	Min	Max
Treated	7,595	0.495	0.500	0	1
Post	7,595	0.217	0.412	0	1
Log(Disintermediation_Score)	7,595	0.235	0.463	0	3.466
Log(Total_Charge)	7,595	4.491	1.617	0.207	11.035
Log(Hours)	2,125	2.528	0.685	0.693	5.537
Communication_Med	7,595	0.364	0.481	0	1
Communication_High	7,595	0.367	0.482	0	1
Client_Personal	7,595	0.532	0.499	0	1

Notes: The unit of analysis is each job assignment. Number of observations in this main analysis sample is 7,595, except for *Log(Hours)*, which is the logarithm of the number of hours the freelancer worked on the assignment, and has values only for 2,125 hourly jobs. *Log(Disintermediation_Score)* is the logarithm of the disintermediation score computed from all messages associated with the assignment plus 1. *Log(Total_Charge)* is the logarithm of the total amount of money charged at the end of the assignment plus 1.

Table 21 compares the three dependent variables for the treated and matched group assignments before and after the Skype block. Before the block, assignments from the two groups are similar in total charges, numbers of hours, and disintermediation scores. After the block, however, the treated group assignments demonstrate significantly more hours, higher total charges, and higher disintermediation scores than the matched group assignments, all suggesting a smaller likelihood of disintermediation.

Table 21: Comparing Treated and Matched Group Job Assignment Outcomes, Before vs. After the Block

Variables	Treatment		Control		Paired t-test
	Mean	Standard error	Mean	Standard error	t-stats
Assignments Before Block:					
Log(Total_Charge)	4.527	0.030	4.473	0.030	-1.276
Log(Hours)	2.557	0.023	2.506	0.024	-1.563
Log(Disintermediation_Score)	0.227	0.009	0.208	0.008	-1.616
Assignments After Block:					
Log(Total_Charge)	4.623	0.054	4.283	0.054	-4.466***
Log(Hours)	2.581	0.044	2.371	0.050	-2.865***
Log(Disintermediation_Score)	0.210	0.015	0.393	0.019	7.502***

Notes: The unit of analysis is each job assignment for the treatment/control group freelancers in the main sample. As in the main analysis, "before the shock" is from Jan to Sep 2017, and "after the shock" is from Feb to Jun 2018. Before the shock, none of the paired t-test results are significant; after the shock, all the paired t-test results are significant. ** significant at 5%; *** significant at 1%.

III.4. Empirical Results

Based on a difference-in-differences estimation framework, the main analyses show that restricting alternative communication technologies outside the platform leads to a decrease in disintermediation on the focal platform. Also, there is a significant increase in the total charge per job and the number of hours per job, suggesting that freelancers in mainland China are less likely to take transactions off the platform after the Skype block compared to unaffected freelancers. The magnitude of the main effect is larger for communication-intensive jobs and for jobs posted by personal clients instead of business clients.

III.4.1. Evidence of Reduced Disintermediation

When investigating how technology affects users' disintermediation tendencies in the marketplace, I find a significant decrease in disintermediation scores and increases in job total charges and job hours after Skype is blocked, all indicating a reduced tendency for disintermediation on the focal platform.

Specifically, the following difference-in-differences regression framework is used to examine how the shock changes the three dependent variables: *Disintermediation_Score*, *Total_Charge*, and *Hours*:

$$Y = \beta_0 + \beta_1 Treated + \beta_2 Post + \beta_3 Treated \times Post + \varepsilon. \tag{5}$$

The unit of analysis is each job assignment. The two groups' job average difference before the shock equals β_1, whereas their average difference after the shock equals $\beta_1 + \beta_3$. Therefore, the estimated coefficient for the interaction term, β_3, captures the shock's impact on the two groups' differences, which is the coefficient of interest here.

Table 22 reports the regression results. Models (1), (3), and (5) use the logarithm of the disintermediation score, the job total charge, and the number of job hours as the dependent variables, respectively. We can see from Models (1), (3), and (5) that, on average, the blocking of Skype seems to have a negative impact on the level of disintermediation on the platform. The estimates of β_3 show that, after Skype is blocked, assignments between mainland China

Table 22: Difference-in-Differences Regressions on the Effect of Skype Block on Disintermediation

Model	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Log(Disintermediation_Score)	Log(Disintermediation_Score)	Log(Total_Charge)	Log(Total_Charge)	Log(Hours)	Log(Hours)
Treated	0.019	0.018***	0.018	0.017	-0.048	-0.049
	[0.011]	[0.004]	[0.042]	[0.063]	[0.034]	[0.067]
Post	0.185***	0.171***	-0.226***	-0.231***	-0.233***	-0.248**
	[0.021]	[0.015]	[0.062]	[0.072]	[0.056]	[0.088]
Treated x Post	-0.202***	-0.204**	0.322***	0.319***	0.257***	0.246**
	[0.027]	[0.014]	[0.087]	[0.072]	[0.074]	[0.081]
Observations	7,595	7,595	7,595	7,595	2,125	2,125
R-squared	0.014	0.014	0.003	0.003	0.006	0.006
Month FE	No	Yes	No	Yes	No	Yes

Notes: Observations are all the job assignments in the main sample. Column (5) and (6) contains only hourly jobs. *Treated* equals 1 if the freelancer is in mainland China. *Post* equals 1 for assignments completed after the block. *Log(Disintermediation_Score)* is the logarithm of the disintermediation score computed from all messages associated with the assignment plus 1. *Log(Total_Charge)* is the logarithm of the total amount of money charged at the end of the assignment plus 1. *Log(Hours)* is the logarithm of the number of hours the freelancer worked on the assignment. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

freelancers and US clients become 20.4% less in detected disintermediation scores, 24.6% higher total charges, and 31.9% longer hours relative to assignments with matched freelancers in Hong Kong, Taiwan, Macau, Singapore, and Malaysia. All three sets of results suggest that the Skype block reduces users' tendencies to disintermediate in the affected regions.

Models (2), (4), and (6) repeat the analysis, controlling for month fixed effects. Adding month fixed effects controls for potential seasonal impacts on the absolute level, allowing the analysis to focus on the changes in the assignment outcomes before and after the block. The main findings continue to hold.

III.4.2. Heterogeneous Effect of Technology on Disintermediation

Various factors moderate the impact of technology on disintermediation; in particular, the effect has a greater magnitude when the job is communication-intensive and when the job is posted by a personal client instead of a business client.

A. Communication-intensive Jobs: Since certain job categories rely more on video or audio chatting to monitor job quality and align the client and the freelancer's expectations, such jobs may require more facilitation during the working process. Hence, the Skype block may have a greater impact when users try to disintermediate on communication-intensive jobs.

Such heterogeneity is tested using a triple difference model that adds a moderator to Equation (1). In Table 23, Model (1) reports the regression results. The negative coefficients of the two three-way interaction terms suggest that jobs that require more facilitation and monitoring during the working process are more significantly affected by the Skype block, with a greater level of decrease in disintermediation tendencies for the most communication-intensive job categories.

These results indicate that blocking Skype hinders the treated users' communication outside the platform and enhances the platform's value creation in facilitating the transactions.

Table 23: Heterogeneity in Disintermediation Tendencies

Model	(1)	(2)
Dependent Variable	Log(Disintermediation_Score)	
Treated	-0.027 [0.023]	0.018 [0.017]
Post	0.111*** [0.043]	0.138*** [0.029]
Treated x Post	-0.103** [0.052]	-0.128*** [0.040]
Communication_Med	-0.011 [0.024]	
Communication_High	-0.020 [0.023]	
Treated x Communication_Med	0.069** [0.032]	
Treated x Communication_High	0.068** [0.030]	
Post x Communication_Med	0.073 [0.052]	
Post x Communication_High	0.130** [0.060]	
Treated x Post x Communication_Med	-0.131** [0.066]	
Treated x Post x Communication_High	-0.151** [0.074]	
Client_Personal		-0.004 [0.016]
Treated x Client_Personal		0.002 [0.023]
Post x Client_Personal		0.091** [0.042]
Treated x Post x Client_Personal		-0.138** [0.055]
Observations	7,595	7,595
R-squared	0.016	0.015

Notes: Observations are all the job assignments in the main analysis sample. *Communication_High* and *Communication_Med* denotes job categories for which the number of milestones is in the top 25%, or the 25% to 75%, of the distribution, respectively. *Client_Personal* equals 1 if the client account type is personal instead of enterprise. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

B. Client Type: The tendency of disintermediation may also vary depending on whether the job is for a business client or a personal client. One potential reason is that enterprise employers are more likely to have their own communication tools or websites for freelancers to contact directly; thus, it is reasonable to assume they generally rely less on Skype for disintermediation relative to personal clients.

Model (2) in Table 23 provides the estimates of the heterogeneity based on whether the job is posted by a personal versus a business client. As expected, the coefficient for the three-way interaction term is -0.30, meaning the Skype block reduces the disintermediation of personal client jobs significantly more than that of enterprise client jobs. Therefore, jobs from personal clients are affected more by the Skype block than business clients.

III.5. Mechanism and Robustness Checks

From where do the above changes come? To understand the mechanism behind the changes in disintermediation tendencies, I find that the affected freelancers are more likely to continue working on the platform after the block compared to the matched freelancers, especially those who used to work on a higher percentage of hourly jobs. As a result, there are a greater number of jobs and a greater percentage of hourly jobs completed by the affected freelancers compared to the matched group after the block. I do not find evidence of selection on job complexity, freelancer qualities, or client qualities. Such increases in the percentage and number of hourly jobs in the treated group is potentially due to the greater cost of disintermediation after Skype is blocked, which forces the freelancers to bring those jobs that would have been disintermediated back to the platform.

III.5.1. Changes in Freelancer Qualities

First, one might argue that selection on freelancer qualities for those who remain active after the block could be the reason for job outcome differences between the two groups. In sum, I do find a compositional change among the remaining freelancers in terms of their past percentage of hourly jobs completed, but no selection on their qualities or working habits.

Summary statistics of the number of freelancers who stayed on the platform in each group show that a greater number of freelancers in the treatment group are more likely to continue working on the platform after the shock compared to the control group. Specifically, 313 freelancers from the treatment group continue to complete jobs on the platform out of the 862 pairs, whereas the number of remaining freelancers is only 263 in the control group.

Are there differences among the remaining freelancers besides the percentage of hourly jobs they completed in the past? Recall that the matching procedure in Section 3 ensured that the freelancers are comparable in 1) their work habits, reflected in the average charge per job, average hours per job, and disintermediation scores and 2) their characteristics, such as job success ratings, availability, and skills. However, since the Skype block introduces additional friction in communication, the clients may tend to hire freelancers who are more capable, more available, or of higher quality to reduce uncertainty in delivering the work, as well as reduce communication costs.

Hence, I test for the potential selection of freelancers after the shock based on their characteristics and past working habits. The t-test results in Panel A of Table 24 show that, after the shock, the freelancers who continue to work in the two groups still have comparable profiles in terms of availability hours, language skills, and total number of tests passed, and there is no significant difference in the remaining freelancers' past job success ratings, ruling out the potential explanation that the freelancers were selected based on their characteristics.

Table 24: Comparing the Treated and Matched Freelancers for Those Who Remained After the Block

Panel A – Comparing the Freelancers' Characteristics

Outcome variable	Treated		Matched		Paired t-test
	Mean	Standard error	Mean	Standard error	t-stats
Freelancer Success Score	0.89	0.01	0.87	0.01	-1.59
Profile Available Hours	39.57	0.14	39.24	0.25	-1.21
Profile Complete Percentage	92.41	1.02	92.49	1.18	0.05
English Level	4.88	0.03	4.91	0.02	0.78
Number of Tests Passed	2.69	0.17	3.06	0.19	1.45

Panel B – Comparing the Freelancers' Past Working Habits

Outcome variable	Treated		Matched		Paired t-test
	Mean	Standard error	Mean	Standard error	t-stats
Past Avg. Job Hours	2.52	0.04	2.44	0.05	-1.38
Past Avg. Job Total Charge	4.53	0.07	4.46	0.08	-0.61
Past Avg. Disintermediation Score	0.23	0.02	0.24	0.02	0.31
Past Percentage of Hourly Jobs	0.295	0.021	0.247	0.021	-1.61

Notes: The unit of analysis is each freelancer in the treated and control group who completed at least one job after Oct 2017. Past job outcome variables are calculated using assignments before the Skype block in Oct 2017, and freelancer characteristics are obtained in Oct 2017 as well. The number of observations is 313 in the treated group and 263 in the matched group, except for: *Past Avg. Job Hours*, which is only available for freelancers who have completed hourly jobs and has 168 observations in the treated group and 131 observations in the matched group; *Freelancer Success Score*, which is only available for freelancers with 3 or more jobs and has 300 observations in the treated group and 240 observations in the matched group. None of the above paired t-test results is significant.

In addition, when comparing the freelancers' past working habits in Panel B of Table 24, we can see that the remaining freelancers again show comparable outcomes; their past jobs have comparable average charges per job and similar disintermediation scores, and the past average number of hours per job is slightly higher for the affected freelancers but not significantly different from the matched group. We do find that the average past percentage of hourly jobs completed is significantly higher among the remaining treatment group freelancers than the control group freelancers. This is consistent with our observation earlier in this section that the freelancers who used to work on more hourly jobs are more likely to stay on the platform after the shock.

In sum, I find that the blocked freelancers are *more* likely to continue completing jobs on the platform after the shock compared to unaffected freelancers, especially those who used to work on a greater percentage of hourly jobs. Apart from this, the freelancers are not significantly different in their 1) job success, availability, and skills or 2) past working habits. Therefore, there is a compositional change but no selection on qualities among the remaining freelancers.

To correct for potential changes in freelancer composition between the two groups, I repeat the main analysis conditional on the freelancers staying on the platform after the shock, which is reported in Table 25. We can see that, after controlling for the freelancer composition change, all main findings still hold.

Table 25: Difference-in-Differences Regressions on the Effect of Skype Block on Disintermediation, Conditional on Staying

Model	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Log(Disintermediation _Score)	Log(Disintermediation _Score)	Log(Total_Charge)	Log(Total_Charge)	Log(Hours)	Log(Hours)
Treated	0.036**	0.036**	0.134**	0.133*	0.045	0.045
	[0.017]	[0.015]	[0.055]	[0.071]	[0.054]	[0.081]
Post	0.187***	0.161***	0.057	0.070	-0.105	-0.132
	[0.023]	[0.018]	[0.066]	[0.054]	[0.067]	[0.086]
Treated x Post	-0.219***	-0.222**	0.206**	0.203**	0.164*	0.164*
	[0.030]	[0.023]	[0.094]	[0.065]	[0.086]	[0.088]
Observations	4,478	4,478	4,478	4,478	1,084	1,084
R-squared	0.020	0.019	0.009	0.009	0.009	0.009
Month FE	No	Yes	No	Yes	No	Yes

Notes: Observations are the job assignments for freelancers who continued working on platform after the shock. Column (5) and (6) contains only hourly jobs. All variables are defined the same as in Table 22. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

III.5.2. Changes in Job Qualities

Given the compositional change of freelancers after the shock, a natural follow-up question is how their jobs become different. I find that, in terms of the total number of jobs, there are slightly more jobs completed by the treatment group than by the control group after Skype is blocked; furthermore, the percentage of hourly jobs becomes significantly higher for the treatment group.

Figure 15 compares the average number of jobs per month that the treated and control freelancers complete before and after the Skype block. While the two groups do not demonstrate any different pre-trends in their numbers of jobs, after the shock, the number of treatment group jobs slowly becomes greater than that of the control group over time. As a result, the total number of jobs of the treatment group (854) is larger than that of the control group (797) after the shock. I observed that the treated freelancers are more likely to stay and continue working on the platform after the shock, which is consistent with the statistics here that they complete more jobs after the shock than the control group freelancers.

Figure 15: Difference between the Treated and Matched Freelancer's Monthly Number of Jobs over Time

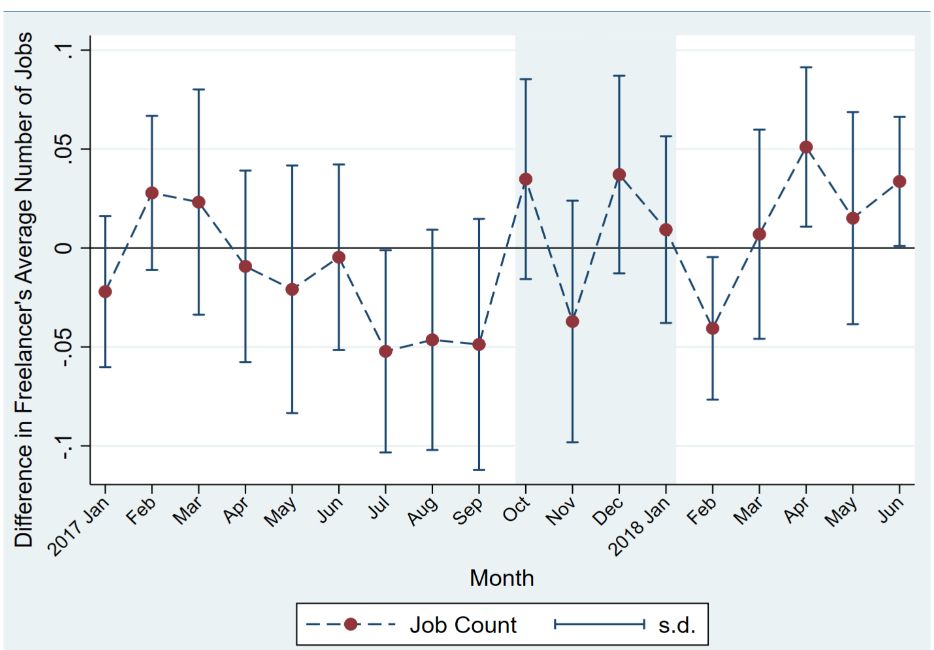

Next, I compare the changes in job type (i.e., whether the jobs are hourly or fixed-priced jobs). Two-sample t-tests show that, while before the shock, there is no significant difference in the percentages of hourly jobs between the treatment group and the control group, after the shock, the percentage is significantly higher ($p = 0.0001$) for the treatment group jobs (0.36) than for the control group jobs (0.18). This difference is again consistent with the compositional change among freelancers; since treated freelancers, who used to work on a greater percentage of hourly jobs, are more likely to continue working after the shock, the proportion of hourly jobs in the treatment group would become higher overall if these treated freelancers continue to prefer hourly jobs.

III.5.3. Changes in Client Qualities

Finally, in terms of the demand side, it is useful to test for any potential changes among the clients who hired the freelancers in each group, since the clients who hired the blocked freelancers might have less intention to disintermediate in the first place. The comparison results suggest that the clients are, in fact, highly comparable between the two groups, indicating that the Skype block did not affect demand on the platform, at least in the short run.

Table 26 reports the t-tests results comparing various characteristics of the clients who hire each group's freelancers before and after the shock. The pre-shock comparison, as expected, shows that there is no difference in the characteristics of the clients who hired the matched pairs of freelancers. After the shock, we can see that the t-test differences are still insignificant; the post-shock clients have similar quality scores, days since registering, percentage of large clients, and percentage of suspended accounts. This suggests that the clients who hire mainland China freelancers after the Skype block remain comparable to those who hire freelancers in the unaffected regions. This, therefore, rules out the possibility of selection bias among clients.

Table 26: Comparing the Clients Who Hired Treated and Control Freelancers, Before and After the Block

Variables	Treated		Control		Paired t-test
	Mean	Standard error	Mean	Standard error	t-stats
			Clients Before the Block		
Client Days Since Register	1251.26	23.87	1288.27	23.91	-1.09
Client Quality Score	65.29	1.27	62.62	1.44	-1.40
% of Large Clients	0.20	0.01	0.19	0.01	0.70
% of Suspended Acct	0.09	0.01	0.08	0.01	-1.57
			Clients After the Block		
Client Days Since Register	1057.98	44.60	1110.11	48.94	-0.79
Client Quality Score	64.70	2.44	63.26	2.82	-0.39
% of Large Clients	0.16	0.02	0.17	0.02	0.41
% of Suspended Acct	0.064	0.011	0.058	0.011	-0.42

Notes: The unit of analysis is each client who hired the treated and control group freelancers in the main sample. None of the above paired t-test results is significant.

A potential reason that we do not observe significant changes on the demand side after the Skype block might be that the analysis sample focuses on short-term changes, and the tasks the clients need are not likely to drastically change in a short period of time. In addition, since the clients tend to be experienced in the hiring process, their job requirements are not likely to change due to technological reasons.

III.5.4. Are the Findings a Result of Lower Efficiency?

Lastly, to corroborate the main findings, I test for the alternative explanation that the results are driven by lowered efficiency from mainland China freelancers when they work without Skype. One might be concerned that the longer working hours and increased total charges are driven by a decrease in the affected freelancers' working efficiency instead of disintermediation.

With lower efficiency, freelancers may work longer hours. However, if the affected freelancers demonstrate lower efficiency, we should not observe an increase in the total charge for fixed-price jobs, as the charge is pre-negotiated. In addition, the results on the disintermediation score also provide more direct evidence of reduced disintermediation regardless of efficiency.

Table 27 repeats the main analysis by job type. We can see that, when examining fixed-price jobs only, there is still a significant increase in the treatment group job's total charge after the shock, meaning that the increase in the charge is the result of less disintermediation instead of lower efficiency.

Table 27: Robustness Check of Treatment Effect on Disintermediation Tendencies, by Job Types

	(1)	(2)	(3)	(4)
Sample	Fixed-Price Jobs		Hourly Jobs	
Dependent Variable	Log(Total_Charge)	Log(Disintermediation_Score)	Log(Total_Charge)	Log(Disintermediation_Score)
Treated	0.146 [0.079]	0.0126 [0.009]	-0.393*** [0.084]	0.031 [0.021]
Post	-0.042 [0.059]	0.160*** [0.017]	-0.398** [0.143]	0.227*** [0.056]
Treated_x_Post	0.182** [0.077]	-0.184*** [0.019]	0.261* [0.122]	-0.283*** [0.055]
Observations	5,470	5,470	2,125	2,125
R-squared	0.005	0.012	0.014	0.014
Month FE	Yes	Yes	Yes	Yes

Notes: The sample is the same as the main analysis sample, split by the job type. All variables are defined the same as in Table 4. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

III.6. Conclusion

While intermediaries focus on understanding internal factors to combat disintermediation, this study shows that external factors, such as the advancement of technology, may also pose an increasing threat of disintermediation to intermediaries and *vice versa*. The study finds that, when Skype is suddenly blocked in mainland China, a major US freelance marketplace has a significant decrease in mainland China freelancers' tendency to disintermediate compared to freelancers in Taiwan, Hong Kong, Macau, Singapore, and Malaysia. The affected freelancers are more likely to continue working on the platform after the block compared to the unaffected freelancers, especially those who used to work on a higher percentage of hourly jobs, potentially due to the increased cost of disintermediation. Therefore, a higher number of jobs and a greater percentage of hourly jobs are completed by the affected freelancers compared to others after the block.

This study aims to open a new research stream for disintermediation studies in online marketplaces, together with prior research, such as the work of Gu and Zhu (2019) in Chapter II. While prior literature mainly focuses on "intermediation," this study shows that "disintermediation" has become the main source of revenue loss, especially with the rapid development of technology. In terms of managerial implications, platform owners should invest in improving their own communication technologies to reduce user disintermediation. Future research can explore the long-term participation and welfare implications of technological change on disintermediation.

APPENDICES

Appendix for "Ideology and Composition among an Online Crowd: Evidence from Wikipedians"

Table A1: Logit Regressions on the Relationship between Contributor Category and Prior Article Category

Model	(1)		(2)		(3)	
Dependent Variable	Contributor Category=-1	Contributor Category=1	Contributor Category=-1	Contributor Category=1	Contributor Category=-1	Contributor Category=1
Prior Article Category	2.0385*** [0.0266]	-2.4427*** [0.0136]	2.0544*** [0.0271]	-2.3705*** [0.0133]	2.0816*** [0.0271]	-2.3158*** [0.0132]
Log(Prior Article Length)			-0.0634*** [0.0045]	0.1200*** [0.0052]	-0.0385*** [0.0051]	0.1590*** [0.0059]
Log(Prior Refs)			-0.2173*** [0.0032]	-0.3061*** [0.0030]	-0.2911*** [0.0043]	-0.4118*** [0.0040]
Year FE	No		No		Yes	
Observations	9,487,164		9,487,164		9,487,164	
Pseudo R-squared	0.021		0.039		0.045	

Notes: The sample is the same as the main analysis sample in Table 3. *Contributor Category* is the categorical version of *Contributor Slant*, which takes the value of -1, 0, or 1, representing contributors with a slant two standard deviations below mean, in between, and above mean, respectively. *Prior Article Category* is the categorical version of *Prior Article Slant*, which takes the value of -1, 0, or 1 representing articles with a slant two standard deviations below mean, in between, and above mean, respectively. Again, we find that the coefficients for the categorical explanatory variable *Prior Article Category* is negative and significant in all cases, suggesting that the slant category of the next contributor is significantly negatively correlated with the slant category of the prior article. Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

Table A2: Regressions on the Relationship between Percentage of Republican in the Area and Prior Article Slant

Model	(1)	(2)
Dependent Variable	RepPerc	RepPerc
Prior Article Slant	-0.0009** [0.0004]	-0.0010** [0.0004]
Log(Prior Article Length)		0.0037*** [0.0001]
Log(Prior Refs)		0.0005*** [0.0001]
Observations	2,438,628	2,438,628
Adjusted R-squared	0.000	0.001

Notes: Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%.

Table A3: Relationship between Contributor Slant and Prior Article Slant, First Edits Only

Models	(1)	(2)
Dependent Variable	Contributor Slant	Contributor Slant
Prior Article Slant	-0.0101***	-0.0238***
	[0.0001]	[0.0005]
Log(Prior Article Length)	0.0007***	0.0011***
	[0.0000]	[0.0001]
Log(Prior Refs)	-0.0005***	-0.0012***
	[0.0000]	[0.0001]
Observations	6,506,072	6,506,072
R-squared	0.008	0.008
Year FE	No	Yes
Article FE	No	Yes
Number of Articles	65,046	65,046

Notes: Robust standard errors in brackets. *significant at 10%; ** significant at 5%; *** significant at 1%. Observations in this panel only include every contributor's first edit of an article.

A4: Additional Robustness Checks on EC vs. Non-EC Effect

We also conduct several additional robustness checks to make sure the Non-EC effect is not driven by alternative explanations. First, our slant index is measured on the basis of frequently used phrases, or code phrases, favored by party representatives. It may be the case that longer articles tend to contain more code phrases and are therefore more measurable. In this case, long articles could drive our results. To rule out this explanation, we eliminate outlying long articles from our full sample, that is, articles that are more than two standard deviations above the mean article length. We obtain similar results.

Second, the articles whose titles contain code phrases might tend to show greater biases in our sample simply because these code phrases are more likely to be used repetitively in the article content. To check our findings against this concern, we exclude from our sample all articles whose title contains code phrases, which is 1.77% of all articles. Again, we find a significant Non-EC effect from the results.

Third, it is possible that certain code phrases are chosen simply because these words do not have other commonly-used synonyms that are neutral or of the opposite slant. In this case, as our measure captures the contributor's choice of words describing the same concept for a given topic, one's contribution may be slanted merely because he or she could not find neutral substitutes of the code phrases to choose from. We rely on the experiences of a legal and copyediting professional to identify these instances in our dictionary and leave only code phrases with natural substitutes. After re-measuring the slant index for articles and contributors, we repeat our analyses and find no significant change in our results. Therefore, the Non-EC effect is not driven by instances where contributors do not have a choice for substitute phrases.

Fourth, because contributors' edits to popular articles tend to have greater impact than those to less popular ones, their political slants measured from these popular articles could carry more weight. Therefore, we use articles' page views as weights when computing the average contribution slant and repeat our analysis using the weighted contributor slant. We continue to find significant Non-EC patterns.

We are also concerned that contributors blocked by Wikipedia administrators may affect our results.[59] These contributors may create extremely biased content initially and drop out of the dataset after being blocked. As a result, contributors overall may become more neutral over time. This problem is mitigated by our approach of assigning missing values to *Contributor Yearly Slant* when a contributor makes no edits in a year. As a robustness check, we repeat our analysis after dropping all 56,329 contributors who have ever been blocked (temporarily or permanently) and the associated 480,960 edits from our sample. Again, the results remain unchanged.

Finally, we test if the Non-EC effect is driven only by extremely slanted articles. We eliminate from our full sample articles with slant index two standard deviation points away from the mean. Changing this threshold to articles without slant in the top and bottom 10% does not differ qualitatively in results. The estimated coefficients with subsamples have the same signs but larger absolute values. We also conduct a robustness check that includes only contributors whose slant is not zero, and we continue to observe a Non-EC pattern among them.

[59] Blocks are used to prevent damage or disruption to Wikipedia. Contributors may be blocked for reasons such as vandalism and edit warring. See https://en.wikipedia.org/wiki/Wikipedia:Blocking_policy for the detailed policy, accessed August 2017.

A5: Procedure for Computing Slant Index

In G&S, for each congressperson c, they observe its ideology y_c and phrase frequency f_{pc}, the number of times phrase p appears in congressperson c's speech, for each phrase p. For each phrase p, G&S regress the relative frequency $\overline{f_{pc}}$, where $\overline{f_{pc}} = f_{pc} / \sum_{p \in P} f_{pc}$, on y_c, and obtain the intercept and slope parameters a_p and b_p, for each phrase p.[60]

The 1,000 phrases exhibit heterogeneous slant. To mitigate the effect of outlier phrases (e.g., "African American" and "illegal immigration"), we set the parameter values for the 9 most left-leaning phrases and 9 most right-leaning phrases to be the same as the 10th most left-leaning phrase and the 10th most right-leaning phrase, respectively.

For each Wikipedia article n, we regress $\overline{f_{pn}}$ - a_p, where $\overline{f_{pn}}$ is the relative frequency of phrase p in the article, on b_p for the 1,000 phrases to obtain the slope estimate $\overline{\mathrm{Y}_\mathrm{n}} = \frac{\sum_{p \in P} b_p (\overline{f_{pn}} - a_p)}{\sum_{p \in P} b_p^2}$.

When an article has none of the 1,000 phrases, $\overline{\mathrm{Y}_\mathrm{n}}$ is 0.4975. We denote $\mathrm{Y}_\mathrm{n} = \overline{\mathrm{Y}_\mathrm{n}} - 0.4975$ and use Y_n as our bias index for article n.

60 The parameter values, together with the 1,000 phrases, are available at http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/26242, accessed March 2019.

Appendix for "Trust and Disintermediation: Evidence from an Online Freelance Marketplace"

Table A6: Summary Statistics and Correlations, Moderators

Variables	Obs.	Mean	Std. dev.	Min	Max	(1)	(2)	(3)	(4)
(1) Same_Country	29,690	0.080	0.272	0	1	1			
(2) Divisible_High	29,690	0.110	0.312	0	1	0.004	1		
(3) Divisible_Med	29,690	0.617	0.486	0	1	0.040	-0.445	1	
(4) Long_Term	12,118	0.405	0.491	0	1	-0.006	-0.194	-0.020	1
(5) Client_Rating_High	29,690	0.534	0.499	0	1	-0.057	-0.076	-0.019	0.195

Notes: The sample of this table is the main analysis sample, the number of observations in this table is 29,690, except for *Long_Term* which includes only jobs with information on their expected duration.

Table A7: Regression of High SSs on Disintermediation with Discrepant vs. Consistent SS and Five-Star Rating

Panel A – High SS, High Five Star Rating

Model	(1)	(2)	(3)
Dependent Variable	Log(Hours)	Log(Total_Charge)	Log(Disintermediation_Score)
Treated	-0.331***	-0.166**	0.112***
	[0.118]	[0.081]	[0.037]
Observations	1,299	2,703	2,199
R-squared	0.006	0.002	0.004

Panel B – High SS, Low Five Star Rating

Model	(1)	(2)	(3)
Dependent Variable	Log(Hours)	Log(Total_Charge)	Log(Disintermediation_Score)
Treated	-0.094*	-0.068*	0.066***
	[0.057]	[0.037]	[0.019]
Observations	4,270	9,512	8,435
R-squared	0.001	0.000	0.001

Notes: The sample in this table is the assignment sample, split into cases where the freelancer's five-star rating is above the 90th percentile. Robust standard errors in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table A8: Regression of High SSs on Disintermediation, Without Overlapping Freelancers

Model	(1)	(2)	(3)
Dependent Variable	Log(Hours)	Log(Total_Charge)	Log(Disintermediation_ Score)
Treated	-0.028	-0.020	0.026*
	[0.040]	[0.027]	[0.014]
SS_High	0.504***	0.541***	-0.141***
	[0.052]	[0.034]	[0.017]
Treated x SS_High	-0.242***	-0.134***	0.063***
	[0.069]	[0.046]	[0.023]
Observations	12,606	28,003	24,427
R-squared	0.011	0.016	0.005

Notes: The analysis sample is the assignment sample for our main analysis, after dropping freelancers who are matched with control group clients after working with treated clients. Robust standard errors in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table A9: OLS Regressions on the Treatment Effect of High SSs on Disintermediation, Using Pre-treatment Data

Model	(1)	(2)	(3)	(4)
Dependent Variable	Log(Past_Hours)	Log(Past_Hours)	Log(Past_Total_Charge)	Log(Past_Total_Charge)
Treated	0.012	0.013	-0.007	0.006
	[0.020]	[0.022]	[0.017]	[0.020]
SS_High	0.404***	0.406***	0.607***	0.624***
	[0.022]	[0.032]	[0.019]	[0.026]
Treated x SS_High		-0.005		-0.035
		[0.044]		[0.037]
Observations	24,848	24,848	31,825	31,825
R-squared	0.015	0.015	0.035	0.035

Notes: Observations are the job assignments in the main analysis with job outcomes replaced by each freelancer's average past job outcome in the 6 months before the study. There are 33,561 assignments from the study; 31,825 involve a freelancer with at least one previous job in the past 6 months and 24,848 involve a freelancer who worked on an hourly job in the past 6 months. *SS_High* and *Treated* are defined as in Table 17. *Log(Past_Hours)* is the logarithm of the average number of hours the freelancer worked on each assignment in the past 6 months. *Log(Past_Total_Charge)* is the logarithm of the freelancer's average total charge per assignment in the past 6 months plus 1. Robust standard errors in brackets. *** p<0.01.

.

Table A10: Repeating the OLS Regressions on the Treatment Effect of High SSs on Disintermediation using a More Robust Disintermediation Score

Model	(1)	(2)
Dependent Variable	Log(Disintermediation_Score_Robust)	Log(Disintermediation_Score_Robust)
Treated	-0.002 [0.002]	-0.005 [0.003]
SS_High	-0.005** [0.002]	-0.009*** [0.003]
Treated x SS_High		0.008** [0.004]
Observations	29,690	29,690
R-squared	0.0002	0.0003

Notes: Observations are all the job assignments created during the study period. The sample and independent variable definitions are the same as in Table 17. *Log(Disintermediation_Score_Robust)* is the logarithm of the robust disintermediation score computed based on the subset of the more explicit keywords plus 1. Robust standard errors in brackets. *** p<0.01, ** p<0.05.

Table A11: First-Time Hires vs. Repeated Hires

Panel A: Regression of High SSs on Disintermediation, First-Time Hires Only

Model Dependent Variable	(1) Log(Hours)	(2) Log(Total_Charge)	(3) Log(Disintermediation_Score)
Treated	-0.0003	-0.018	0.027*
	[0.040]	[0.027]	[0.014]
SS_High	0.439***	0.578***	-0.087***
	[0.050]	[0.035]	[0.017]
Treated x SS_High	-0.124*	-0.084*	0.047**
	[0.069]	[0.048]	[0.023]
Observations	12,694	26,469	23,577
R-squared	0.010	0.020	0.002

Panel B: Regression of High SSs on Disintermediation, Repeated Hires Only

Model Dependent Variable	(1) Log(Hours)	(2) Log(Total_Charge)	(3) Log(Disintermediation_Score)
Treated	-0.024	0.059	-0.018
	[0.099]	[0.049]	[0.027]
SS_High	0.318***	0.302***	-0.265***
	[0.121]	[0.056]	[0.030]
Treated x SS_High	-0.282*	-0.138*	0.096**
	[0.166]	[0.079]	[0.044]
Observations	1,899	7,092	6,113
R-squared	0.005	0.006	0.017

Notes: The sample in the above table is the main analysis sample, with Model (1) in both panels including only hourly jobs, and Models (2) and (3) including both hourly and fixed-price jobs. Model (3) contains jobs for which disintermediation score is available so the number of observations is smaller than Model (2). Robust standard errors in brackets. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A12: OLS Regressions on the Treatment Effect of High SSs on

Client Feedback of Freelancers

Model	(1)	(2)
Dependent Variable	Client_Feedback	Client_Feedback
Treated	-0.001	-0.006
	[0.008]	[0.011]
SS_High	0.102***	0.094***
	[0.007]	[0.011]
Treated x SS_High		0.016
		[0.015]
Observations	23,501	23,501
R-squared	0.007	0.007

Notes: Observations are the job assignments in the main analysis with non-missing client feedback. Independent variable definitions are the same as in Table 17. *Client_Feedback* is the five-star feedback rating that the client left when the job assignment was closed. Robust standard errors in brackets. *** p<0.01.

REFERENCES

Adner, R., and Kapoor, R. 2010. Value creation in innovation ecosystems: How the structure of technological interdependence affects firm performance in new technology generations. *Strategic Management Journal* 31: 306–333.

Afuah, A., and Tucci, C.L. 2012. Crowdsourcing as a solution to distant search. *Academy of Management Review* 37(3): 355–375.

Agrawal, A., Lacetera, N., and Elizabeth Lyons. 2016. Does standardized information in online markets disproportionately benefit job applicants from less developed countries? *Journal of International Economics* 103(1): 1–12.

Antweiler, W., and Frank, M.Z. 2004. Is all that talk just noise? The information content of internet stock message boards. *Journal of Finance* 59(3): 1259–1294.

Arazy, O., Nov, O., Patterson, R., and Yeo, L. 2011. Information quality in Wikipedia: The effects of group composition and task conflict. *Journal of Management Information Systems* 21(4): 71–98.

Arazy, O., Ortega, F., Nov, O., Yeo, L., and Balila, A. 2015. Functional roles and career paths in Wikipedia. *Computer Supported Cooperative Work (CSCW):* 1092–1105.

Armstrong, Mark. 2006. Competition in two-sided markets. *The RAND Journal of Economics* 37(3): 668–691.

Arya, A., Mittendorf, B., and Sappington, D.E.M. 2007. The bright side of supplier encroachment. *Marketing Science* 26(5): 651–659.

Ba, S., and Pavlou, P.A. 2002. Evidence of the effect of trust building technology in electronic markets: Price premium and buyer behavior. *MIS Quarterly* 26(3): 243–268.

Bail, C.A., Argyle, L.P., Brown, T.W., Bumpus, J.P., Chen, H., Hunzaker, M.F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. 2018. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences* 115(37): 9216–9221.

Bagozzi, R.P., and Dholakia, U.M. 2006. Open source software user communities: A study of participation in Linux user groups. *Management Science* 52(7): 1099–1115.

Bakshy, E., Messing, S. and Adamic, L.A. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348(6239):1130-1132.

Baldwin, C.Y., and Clark, K.B. 2000. *Design Rules: The Power of Modularity*. Vol. 1. Cambridge, MA: MIT Press.

Bapna, R., Qiu, L., and Rice, S. 2017. Repeated interactions versus social ties: Quantifying the economic value of trust, forgiveness, and reputation using a field experiment. *MIS Quarterly* 41(3): 841–866.

Barach, M., Kaul, A., Leung, M., and Lu, S. 2018. Small numbers bargaining in the age of big data: Evidence from a two-sided labor matching platform. Available at SSRN: https://ssrn.com/abstract=3277455.

Bassamboo, A., Cui, R., and Moreno, A. 2015. The wisdom of crowds in operations: Forecasting using prediction markets. Working paper. SSRN: https://ssrn.com/abstract=2679663.

Benkler, Y. 2006. *The Wealth of Networks: How Social Production Transforms Markets and Freedom.* Yale University Press.

Blake, T., and Coey, D. 2014. Why marketplace experimentation is harder than it seems: The role of test-control interference. *Proceedings of the Fifteenth ACM Conference on Economics and Computation* 567–582.

Bloom, N., Garicano, L., Sadun, R., and Van Reenen, J. 2014. The distinct effects of information technology and communication technology on firm organization. *Management Science* 60(12): 2859–2885.

Bohnet, I., and Huck, S. 2004. Repetition and reputation: Implications for trust and trustworthiness when institutions change. *American Economic Review* 94(2): 362–366.

Bolton, G., Greiner B., and Ockenfels, A. 2013. Engineering trust: reciprocity in the production of reputation information. *Management Science* 59(2): 265–285.

Boudreau, K., and Jeppesen, L. 2015. Unpaid crowd complementors: The platform network effect mirage. *Strategic Management Journal* 36: 1761–1777.

Boxell, L., Gentzkow, M., and Shapiro, J.M. 2017. Is the internet causing political polarization? Evidence from demographics. *National Bureau of Economic Research* Working paper No. 23258.

Budak, C., Goel, S., and Rao, J. M. 2016. Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly* 80(S1): 250–271.

Burke, M., and Kraut, R. 2008. Mopping up: Modeling Wikipedia promotion decisions. *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (CSCW)*: 27–36.

Cai, H., Jin, G.Z., Liu, C., and Zhou, L-A. 2013. More trusting, less trust? An investigation of early e-commerce in China. *National Bureau of Economic Research Working Paper No. w18961*.

Cabral, L., and Hortaçsu, A. 2010. The dynamics of seller reputation: evidence from eBay. *Journal of Industrial Economics* 58(1): 54–78.

Cabral, L., Ozbay, E.Y., and Schotter, A. 2014. Intrinsic and instrumental reciprocity: An experimental study. *Games and Economic Behavior* 87: 100–121.

Carr, N. 2008. *The Big Switch: Rewiring the World, from Edison to Google*. Norton, New York.

Cennamo, C., and Santalo, S. 2013. Platform competition: Strategic trade-offs in platform markets. *Strategic Management Journal* 34: 1331–1350.

Charness, G., Du, N., and Yang, C-L. 2011. Trust and trustworthiness reputations in an investment game. *Games and Economic Behavior* 72: 361–375.

Chen, S.S., Chuang, Y.W., and Chen, P.Y. 2012. Behavioral intention formation in knowledge sharing: Examining the roles of KMS quality, KMS self-efficacy, and organizational climate. *Knowledge-Based Systems* 31(7): 106–118.

Chesbrough, H.W. 2006. *Open Innovation: the New Imperative for Creating and Profiting from Technology.* Harvard Business Press.

Chesney, T. 2006. An empirical examination of Wikipedia's credibility. *First Monday* 11(11).

Chiang, C.F., and Knight, B. 2011. Media bias and influence: Evidence from newspaper endorsements. *Review of Economic Studies* 78(3): 795–820.

Collier, B., Burke, M., Kittur, N., and Kraut, R. 2008. Retrospective versus prospective evidence for promotion: The case of Wikipedia. *2008 Meeting of the Academy of Management.*

Dai, W., Jin, G.Z., Lee, J., and Luca, M. 2018. Aggregation of consumer ratings: An application to Yelp.com. *Quantitative Marketing and Economics* 16(3): 289–339.

Dellarocas, C. 2003. The digitization of word of mouth: promise and challenges of online feedback mechanisms. *Management Science* 49(10): 1407–1424.

DellaVigna, S., and Kaplan, E. 2007. The Fox News effect: Media bias and voting. *Quarterly Journal of Economics* 122(3): 1187–1234.

Draper, N.R., and Smith, H. 1998. *Applied Regression Analysis (3rd ed.).* John Wiley.

Durante, R., and Knight, B. 2012. Partisan control, media bias, and viewer responses: Evidence from Berlusconi's Italy. *Journal of the European Economic Association* 10(3): 451–481.

Edelman, B., and Hu, P. 2016. Disintermediation in two-sided marketplaces. *Harvard Business School Technical Note 917–004*.

Edelman, B., and Wright, J. 2015. Price coherence and excessive intermediation. *Quarterly Journal of Economics* 130(3): 1283–1328.

Einav, L., Farronato, C., and Levin, J. 2016. Peer-to-peer markets. *Annual Review of Economics* 8: 615–635.

Ert, E., Fleischer, A., and Magen, N. 2015. Trust and reputation in the sharing economy: The role of personal photos in Airbnb." in *Advances in Consumer Research.* Vol. 43, eds. Kristin Diehl and Carolyn Yoon, Duluth, MN: Association for Consumer Research, 518–519.

Fan, M., Tan, Y., and Whinston, A.B. 2005. Evaluation and design of online cooperative feedback mechanisms for reputation management. *IEEE Transactions on Knowledge and Data Engineering* 17(2): 244–254.

Faraj, S., Jarvenpaa S.L., and Majchrzak, A. 2011. Knowledge collaboration in online communities. *Organization Science* 22(5): 1224–1239.

Forte, A., Kittur, N., Larco, V., Zhu, H., Bruckman, A., and Kraut, R.E. 2012. Coordination and beyond: Social functions of groups in open content production. *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work:* 417–426.

Forte, A., Larco, V., and Bruckman, A. 2009. Decentralization in Wikipedia governance. *Journal of Management Information Systems* 26(1): 49–72.

Fuentelsaz, L., Garrido, E., and Maicas, JP. 2014. "Incumbents, Technological Change and Institutions: How The Value of Complementary Resources Varies across Markets." *Strategic Management Journal* 36: 1778–1801.

Gallus, J. 2016. Fostering public good contributions with symbolic awards: A large-scale natural field experiment at Wikipedia. *Management Science* 63(12): 3999–4015.

Galton, F. 1907. Vox populi (the wisdom of crowds). *Nature* 75(7): 450–451.

Garimella, K., De Francisci Morales, G., Gionis, A. and Mathioudakis, M. 2018. Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference*.

Gentzkow, M., and Shapiro, J.M. 2003. Media, education, and anti-Americanism in the Muslim world. *Journal of Economic Perspectives* 18(3): 117–133.

Gentzkow, M., and Shapiro, J.M. 2010. What drives media slant? Evidence from U.S. daily newspapers. *Econometrica* 78(1): 35–71.

Gentzkow, M., and Shapiro, J.M. 2011. Ideological segregation online and offline. *Quarterly Journal of Economics* 126(4): 1799–1839.

Gilbert, E., Bergstrom, T., and Karahalios, K. 2009. Blogs are echo chambers: Blogs are echo chambers. In *42nd Hawaii International Conference on System Sciences*.

Giles, J. 2005. Internet encyclopaedias go head to head. *Nature* 438(12): 900–901.

Greenstein, S., and Zhu, F. 2012. Is Wikipedia biased? *American Economic Review (Papers and Proceedings)* 102(3): 343–348.

Greenstein, S., and Zhu, F. 2016. Open content, Linus' Law, and neutral point of view. *Information Systems Research* 27(3): 618–635.

Greenstein, S., and Zhu, F. 2018. Do experts or collective intelligence write with more bias? Evidence from Encyclopædia Britannica and Wikipedia. *MIS Quarterly* 42(3): 945–959.

Gulati, R. 1995. Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances." *Academy of Management Journal* 38(1): 85–112.

Hagiu, A., and Wright, J. 2015. Multi-sided platforms. *International Journal of Industrial Organization* 43: 162–174.

Halfaker, A., Kittur, A., and Riedl, J. 2011. Don't bite the newbies: How reverts affect the quantity and quality of Wikipedia work. *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym'11).*

Hargadon, A.B., and Bechky, B.A. 2006. When collections of creatives become creative collectives: A field study of problem solving at work. *Organization Science* 17(4): 484–500.

Hill, M. 2017. Almost Wikipedia: Eight early Encyclopedia projects and the mechanisms of collective action. Working paper.

Hinds, P., and Kiesler, S. 1995. Communication across boundaries: Work, structure, and use of communication technologies in a large organization. *Organization Science* 6(4): 373–393.

Horton, J., and Golden, J. 2015. Reputation inflation: Evidence from an online labor market. *Working Paper, New York University*.

Horton, J., Rand, D., and Zeckhauser, R. 2011. The online laboratory: Conducting experiments in a real labor market. *Experimental Economics* 14(3): 399–425.

Hu, N., Pavlou, P.A., and Zhang, J. 2009. Overcoming the J-shaped distribution of product reviews. *Communications of the ACM* 52(10).

Hui, X., Saeedi, M., Shen, Z., and Sundaresan, N. 2014. From lemon markets to managed markets: the evolution of eBay's reputation system." *Ohio State University Working Paper*.

Iansiti, M., and Levien, R. 2004. *The Keystone Advantage: What the New Dynamics of Business Ecosystems Mean for Strategy, Innovation, and Sustainability*. Harvard Business School Press: Boston, MA.

Jelveh, Z., Kogut, B., and Naidu, S. 2014. Political language in economics. *Columbia Business School Research,* Paper No. 14–57.

Jemielniak, D. 2014. *Common Knowledge?: An Ethnography of Wikipedia*. Palo Alto, CA: Stanford University Press.

Jeppesen, L.B., and Frederiksen, L. 2006. Why do users contribute to firm-hosted user communities? The case of computer-controlled music instruments. *Organization Science* 17(1): 45–63.

Jin, G.Z., and Kato, A. 2006. Price, quality, and reputation: Evidence from an online field experiment. *The RAND Journal of Economics* 37(4): 983–1005.

Kane, G.C., and Fichman, R.G. 2009. The shoemaker's children: Using Wikis for information systems teaching, research, and publication. *MIS Quarterly* 33(1): 1–17.

Kane, G.C., Johnson, J., and Majchrzak, A. 2014. Emergent life cycle: The tension between knowledge change and knowledge retention in open online coproduction communities. *Management Science* 60(12): 3026–3048.

Kapoor, A., and Tucker, C.E. 2017. How do platform participants respond to an unfair rating? An analysis of a ride-sharing platform using a quasi-experiment. Available at SSRN: https://ssrn.com/abstract=2970772.

Kapoor, R., and Lee, JM. 2013. Coordinating and competing in ecosystems: How organizational forms shape new technology investments. *Strategic Management Journal* 34: 274–296.

Kim, H-W., Xu, Y., and Gupta, S. 2012. Which is more important in internet shopping, perceived price or trust? *Electronic Commerce Research and Applications* 11: 241–252.

Kittur, A., Chi, E., Pendleton, B.A., Suh, B., and Mytkowicz, T. 2007. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. *World Wide Web* 1(2): 19.

Kittur, A., and Kraut, R. 2008. Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. *ACM Conference on Computer Supported Cooperative Work:* 37–46.

Kittur, A., Suh, B., Pendleton, B.A., and Chi, E.H., 2007. He says, she says: Conflict and coordination in Wikipedia. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems:* 453–462.

Kogut, B., and Metiu, A. 2000. The emergence of E-Innovation: Insights from open source software development. *Reginald H. Jones Center Working Paper.*

Kogut, B., and Metiu, A. 2001. Open-source software development and distributed innovation. *Oxford Review of Economic Policy* 17(2): 248–264.

Kogut, B., and Zander, U. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science* 3(3): 383–397.

Kuk, G. 2006. Strategic interaction and knowledge sharing in the KDE developer mailing list. *Management Science* 52(7): 1031–1042.

Lawrence, E., Sides, J., and Farrell, H. 2010. Self-segregation or deliberation?: Blog readership, participation, and polarization in American politics. *Perspectives on Politics* 8(1): 141–157.

Lee, G.K., and Cole, R.E. 2003. From a firm-based to a community-based model of knowledge creation: The case of the Linux kernel development. *Organization Science* 14(6): 633–649.

Lee, Y-J., Hosanagar, K., and Tan, Y. 2015. Do I follow my friends or the crowd? Information cascades in online movie ratings. *Management Science* 61(9): 2241–2258.

Lemos, R. 2004. Security research suggests linux has fewer flaws. *CNET News*, December 13, http://news.cnet.com/Security-research-suggests-Linux-has-fewer-flaws/2100-1002_3-5489804.html.

Lerner, J., and Tirole, J. 2002. Some simple economics of open source. *Journal of Industrial Economics* 50(2): 197–234.

Liu, Q. 2011. Information acquisition and reputation dynamics. *Review of Economic Studies* 78(4): 1400–1425.

Luca, M. 2017. Designing online marketplaces: Trust and reputation mechanisms. *NBER Innovation Policy and the Economy*, 77–93.

Majchrzak, A. 2009. Where is the theory in Wikis? *MIS Quarterly* 33(1): 18–20.

Masterov, D., Meyer, U., and Tadelis, S. 2015. Canary in the E-Commerce coal mine: Detecting and predicting poor experiences using buyer-to-seller messages. *Proceedings of the Sixteenth ACM Conference on Economics and Computation*.

McPherson, M., Smith-Lovin, L., and Cook, J.M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27(1): 415–444.

Mollick, E., and Nanda, R. 2015. Wisdom or madness? Comparing crowds with expert evaluation in funding the arts. *Management Science* 62(6): 1533–1553.

Moon, J.Y., and Sproull, L.S. 2008. The role of feedback in managing the Internet-based volunteer work force. *Information Systems Research* 19(4): 494–515.

Moreno, A., and Terwiesch, C. 2014. Doing business with strangers: Reputation in online service marketplaces. *Information Systems Research* 25(4): 865–886.

Muchnik, L., Aral, S., and Taylor, S.J. 2013. Social influence bias: A randomized experiment. *Science* 341(6146): 647–651.

Nagaraj, A. 2017. Does copyright affect reuse? Evidence from Google Books and Wikipedia. *Management Science* 64(7): 3091–3107.

Nagle, F. 2018. Learning by contributing: Gaining competitive advantage through contribution to crowdsourced public goods. *Organization Science* 29(4): 569–587.

Nambisan, S., and Baron, R.A. 2010. Different roles, different strokes: Organizing virtual customer environments to promote two types of customer contributions. *Organization Science* 21(2): 554–572.

Nyhan, B., and Reifler, J. 2010. When corrections fail: The persistence of political misperceptions. *Political Behavior* 32(2): 303–330.

Park, J., Konana, P., Gu, B., Kumar, A., and Raghunathan, R. 2013. Information valuation and confirmation bias in virtual communities: Evidence from stock message boards. *Information Systems Research* 24(4): 1050–1067.

Parker, G., and Van Alstyne, MW. 2005. Two-sided network effects: A theory of information product design. *Management Science* 51(10): 1494–1504.

Pavlou, P.A., and Dimoka, A. 2006. The nature and role of feedback text comments in online marketplaces: Implications for trust building, price premiums, and seller differentiation. *Information Systems Research* 17(4): 392–414.

Pavlou, P.A., Liang, H., and Xue, Y. 2007. Understanding and mitigating uncertainty in online exchange relationships: A principal-agent perspective. *MIS Quarterly* 31(1): 105–136.

Peukert, C., and Reimers, I. 2018. Digital disintermediation and efficiency in the market for ideas. *CESifo Working Paper Series No. 6880*.

Pierce, L. 2012. Organizational structure and the limits of knowledge sharing: Incentive conflict and agency in car leasing. *Management Science* 58(6): 1106–1121.

Piskorski, M.J., and Gorbatâi A. 2017. Testing Coleman's social-norm enforcement mechanism: Evidence from Wikipedia. *American Journal of Sociology* 122(4): 1183–1222.

Prelec, D., Seung, H.S. and McCoy, J. 2017. A solution to the single-question crowd wisdom problem. *Nature* 541(7638): 532.

Puranam, P., and Vanneste, B. Trust and governance: Untangling a tangled web. *Academy of Management Review*, forthcoming.

Purdy, W. 2015. Radicalization: Social media and the rise of terrorism. *Testimony presented before the House Committee on Oversight and Government Reform's Subcommittee on National Security*, https://oversight.house.gov/wp-content/uploads/2015/10/10-28-2015-Natl-Security-Subcommittee-Hearing-on-Radicalization-Purdy-TRC-Testimony.pdf.

Qiu, L., and Kumar, S. 2017. Understanding voluntary knowledge provision and content contribution through a social-media-based prediction market: A field experiment. *Information Systems Research* 28(3): 529–546.

Quattrociocchi, W., Scala, A., and Sunstein, C.R. 2016. Echo chambers on Facebook. Working paper. SSRN: https://ssrn.com/abstract=2795110.

Rajagopalan, M. S., Khanna, V. K., Leiter, Y., Stott, M., Showalter, T. N., Dicker, A. P., and Lawrence, Y. R. 2011. Patient-oriented cancer information on the internet: a comparison of Wikipedia and a professionally maintained database. *Journal of Oncology Practice* 7(5): 319–323.

Ramasubbu, N., and Balan, R.K. 2012. Overcoming the challenges in cost estimation for distributed software projects. *Proceedings of the 34th International Conference on Software Engineering*: 91–101.

Ransbotham, S., and Kane, G.C. 2011. Membership turnover and collaboration success in online communities: Explaining rises and falls from grace in Wikipedia. *MIS Quarterly* 35(3): 613–627.

Ransbotham, S., Kane, G.C., and Lurie, N.H. 2012. Network characteristics and the value of collaborative user-generated content. *Marketing Science* 31(3): 387–405.

Rector, L. H. 2008. Comparison of Wikipedia and other encyclopedias for accuracy, breadth, and depth in historical articles. *Reference Services Review* 36(1): 7–22.

Resnick, P., and Zeckhauser, R. 2002. Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. *Advances in Applied Microeconomics* 11: 127–157.

Rider, C., and Samila, A. 2019. Envisioning value: Certification, matchmaking, and returns to brokerage. *Working paper*.

Roberts, J.A., Hann, I.H., and Slaughter, S.A. 2006. Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the Apache projects. *Management Science* 52(7): 984–999.

Rochet, J-C., and Tirole, J. 2006. Two-sided markets: A progress report. *The RAND Journal of Economics* 37(3): 645–667.

Rothaermel, F.T., and Sugiyama, S. 2001. Virtual Internet communities and commercial success: Individual and community-level theory grounded in the atypical case of TimeZone.com. *Journal of Management* 27(3): 297–312.

Schroeder, A., and Wagner, C. 2012. Governance of open content creation: a conceptualization and analysis of control and guiding mechanisms in the open content domain. *Journal of the American Society for Information Science and Technology* 63(10): 1947–1959.

Shankar, V., and Bayus, B. 2003. Network effects and competition: an empirical analysis of the home video game industry. *Strategic Management Journal* 24: 375–384.

Shankland, S. 2003. Study lauds open-source code quality. *CNET News*, February 19. http://news.cnet.com/Study-laudsopen-source-code-quality/2100-1001_3-985221.html.

Shore, J., Baek, J., and Dellarocas, C. 2016. Network structure and patterns of information diversity on Twitter. Working paper.

Sobel, J. 2005. Interdependent preferences and reciprocity. *Journal of Economic Literature* 43(2): 392–436.

Spulber, D.F. 1996. Market microstructure and intermediation. *Journal of Economic Perspectives* 10(3): 135–152.

Spulber, D.F. 1996. Market making by price-setting firms. *Review of Economic Studies* 63(4): 559–580.

Spulber, D.F. 1999. Market microstructure: Intermediaries and the theory of the firm. New York, NY: Cambridge University Press.

Spulber, D.F. 2011. Should business method inventions be patentable? *Journal of Legal Analysis* 3(1): 265–340.

Stanton, C.T., and Thomas, C. 2016. Landing the first job: The value of intermediaries in online hiring. *Review of Economic Studies* 83(2): 810–854.

Stone, D.F. 2011. Ideological media bias. *Journal of Economic Behavior & Organization* 78(3): 256–271.

Strader, T., and Ramaswami, S. 2002. The value of seller trustworthiness in C2C online markets. *Communications of the ACM* 45(12): 45–49.

Stvilia, B., Twidale, M. B., Smith, L. C., and Gasser, L. 2008. Information quality work organization in Wikipedia. *Journal of the American Society for Information Science and Technology* 59(6): 983–1001.

Sun, T., Viswanathan, S., and Zheleva, E. 2017. Creating social contagion through firm mediated message design: Evidence from a randomized field experiment. *Robert H. Smith School Research Paper No. RHS 2543864.*

Sunstein, C.R. 2001. *Echo Chambers: Bush v. Gore, Impeachment, and Beyond.* Princeton University Press.

Surowiecki, J. 2004. The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business. *Economies, Societies and Nations* 296.

Tan, C.H., Sutanto, J., Phang, C.W., and Gasimov, A. 2014. Using personal communication technologies for commercial communications: A cross-country investigation of Email and SMS." *Information Systems Research* 25(2): 307–327.

Te'eni, D. 2009. Comment: The Wiki way in a hurry—The ICIS anecdote. *MIS Quarterly* 33(1): 20–22.

ter Huurne, M., Ronteltap, A., Corten, R., and Buskens, V. 2017. Antecedents of trust in the sharing economy: A systematic review. *Journal of Consumer Behavior* 16(6): 485–498.

Tucci, C.L., and Villarroel, J.A. 2007. Open Source Corporate Strategy (OSCS): Reaching outside the firm's boundaries to tap upon global talent. In *Academy of Management Annual Meeting* (No. CSI-CONF-2007-001).

Utz, S., Matzat, U., and Snijders, C. 2009. On-line reputation systems: the effects of feedback comments and reactions on building and rebuilding trust in on-line auctions. *International Journal of Electronic Commerce* 13(3): 95–118.

Van Alstyne, M., and Brynjolfsson, E. 2005. Global village or cyber-balkans? Modeling and measuring the integration of electronic communities. *Management Science* 51(6): 851–868.

Von Krogh, G., and Von Hippel, E. 2006. The promise of research on open source software. *Management Science* 52(7): 975–983.

Waldfogel, J. 2012. And the bands played on: Digital disintermediation and the quality of new recorded music." Available at SSRN: https://ssrn.com/abstract=2117372.

Waldfogel, J., and Reimers, I. 2015. Storming the gatekeepers: Digital disintermediation in the market for books. *Information Economics and Policy* 31(1): 47–58.

Venkatesh, V., Bala, H. and Sambamurthy, V. 2016. Implementation of an information and communication technology in a developing country: A multimethod longitudinal study in a bank in India. *Information Systems Research* 27(3): 558–579.

Wang, Y., Goes, P., Wei, Z., and Zeng, D.D. 2017. Production of online word-of-mouth: Peer effects and the moderation of user characteristics. Working paper. SSRN: https://ssrn.com/abstract=2899452.

Wasko, M.M., and Faraj, S. 2005. Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. *MIS Quarterly* 29(1): 35–57.

Wood, T., and Porter, E. 2019. The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior* 41(1): 135–163.

Xu, L., Niculescu, M.F., and Wu, D.J. 2015. When to use the open business model for software products under network effects? Working paper.

Xu, S.X., and Zhang, X.M. 2009. How do social media shape the information environment in the financial market? *ICIS 2009 Proceedings*: 56.

Xu, S.X., and Zhang, X.M. 2013. Impact of Wikipedia on market information environment: Evidence on management disclosure and investor reaction. *MIS Quarterly* 37(4): 1043–1068.

Xu, X., Thong, J.Y., and Venkatesh, V. 2014. Effects of ICT service innovation and complementary strategies on brand equity and customer loyalty in a consumer technology market. *Information Systems Research* 25(4): 710–729.

Yang, Y., Chen, P.Y., and Pavlou, P. 2009. Open innovation: An empirical study of online contests. *ICIS 2009 Proceedings*: 13.

Zhang, X.M., and Zhu, F. 2011. Group size and incentives to contribute: A natural experiment at Chinese Wikipedia. *American Economic Review* 101(4): 1601–1615.

Zhu, F., Chen, W., and Sun, S. 2018. ZBJ: Building a global outsourcing platform for knowledge workers (A). *Harvard Business School Case 618-044*.

Zhu, F., and Iansiti, M. 2012. Entry into platform-based markets. *Strategic Management Journal* 33: 88–106.