

# Gender Inequality in New Media: Evidence from Wikipedia\*

Marit Hinnosaar<sup>†</sup>

April 26, 2019

## Abstract

Media is critical for gender equality. I analyze Wikipedia, one of the prominent examples of new media. Using data from a survey and a randomized survey experiment, I study why women are less likely to contribute to Wikipedia, the implications of the gender gap on Wikipedia's content, and what can be done about it. I find that: (1) gender differences in the frequency of Wikipedia use and in beliefs about one's competence explain a large share of the gender gap in Wikipedia writing; (2) the gender gap among contributors leads to unequal coverage of topics; (3) providing information about gender inequality has a large effect on contributions.

**JEL codes:** L86; L82; J16; H41

**Keywords:** Gender, Media, Internet, Public goods

## 1 Introduction

In most traditional media outlets, men vastly outnumber women, and several major media organizations have taken steps to change that. Examples include Bloomberg introducing gender quotas for news stories<sup>1</sup> and BBC for panel shows.<sup>2</sup> At the global level, the United

---

\*I'm grateful to Ainhoa Aparicio, Gabriele Camera, Matthias Doepke, Martin Dufwenberg, Ruben Durante, Avi Goldfarb, Shane Greenstein, Toomas Hinnosaar, Eliana La Ferrara, Ryan McDevitt, David McKenzie, Mario Pagliero, Andrea Prat, Stephan Seiler, Alex Tetenov, and Michael Zhang, as well as seminar audiences at the European University Institute, ZEW Conference on the Economics of ICT, Digital Economy Workshop, ICT Paris Conference, and Economics of Media Bias Workshop in Cologne for valuable comments.

<sup>†</sup>Collegio Carlo Alberto and CEPR, [marit.hinnosaar@gmail.com](mailto:marit.hinnosaar@gmail.com).

<sup>1</sup>Posetti, Julie. 2015. "The Bloomberg News recipe for newsroom transformation." World News Publishing Focus by World Association of Newspapers and News Publishers, March 17. <http://blog.wan-ifra.org/2015/03/17/the-bloomberg-news-recipe-for-newsroom-transformation>.

<sup>2</sup>BBC. 2014. "BBC TV boss: No all-male panel shows." February 9. <http://www.bbc.com/news/entertainment-arts-26107011>.

Nations is engaged in an initiative to increase female presence in the media.<sup>3</sup> One reason to introduce these measures is that historically rigid institutions, work arrangements, and hiring policies in the media might otherwise inhibit women to be heard. In this paper, I analyze a new media source, which is not restricted by these institutions, but which nevertheless faces similar issues.

I analyze Wikipedia, which is a widely used online encyclopedia. It is the fifth most popular website in the world.<sup>4</sup> Although anyone can write and edit Wikipedia articles, there is evidence that the majority of Wikipedia writers are men. The Wikimedia Foundation, which owns Wikipedia, views the lack of female contributors as a problem and has set itself goals to increase their share.<sup>5</sup>

Why does it matter who contributes to Wikipedia? Who contributes to its production matters if a gender gap among contributors leads to biased content. There is some evidence that Wikipedia's coverage of topics is gender biased. In particular, popular media has voiced concerns that Wikipedia has less information about female than male public figures,<sup>6</sup> writers,<sup>7</sup> and scientists.<sup>8</sup> For example, in 2018 a female Nobel Prize winner didn't have a Wikipedia page.<sup>9</sup> More generally, Reagle and Rhue (2011) document that biographies of women are more likely to be missing from Wikipedia compared to *Encyclopedia Britannica* than those of men.

Coverage of public figures in Wikipedia is important because they are role models. Literature provides evidence that role models among other things influence educational decisions and career choices, which affect earnings potential and overall economic outcomes. For example, using a randomized natural experiment in India, Beaman et al. (2012) find that exposure to women in local government eliminates the gender gap in educational attainment. In a field experiment, Porter and Serra (2017) find that exposure to female economics alumni

---

<sup>3</sup><http://www.unesco.org/new/en/communication-and-information/crosscutting-priorities/gender-and-media/global-alliance-on-media-and-gender/about-gamag/>, accessed November 15, 2017.

<sup>4</sup>The popularity of Wikipedia is measured by the web traffic measurement company Alexa Internet (<http://www.alexa.com/siteinfo/wikipedia.org>, accessed November 15, 2017).

<sup>5</sup>BBC. 2014. "Wikipedia 'completely failed' to fix gender imbalance." BBC interview with Jimmy Wales, August 8. <http://www.bbc.com/news/business-28701772>.

<sup>6</sup>Cohen, Noam. 2011. "Define Gender Gap? Look Up Wikipedia's Contributor List." New York Times, January 30. <http://www.nytimes.com/2011/01/31/business/media/31link.html>.

<sup>7</sup>Flood, Alison. 2013. "Wikipedia bumps women from 'American novelists' category." The Guardian, April 25. <https://www.theguardian.com/books/2013/apr/25/wikipedia-women-american-novelists>.

<sup>8</sup>Donald, Athene, and Frank Norman. 2013. "Using Wikipedia to inspire the next generation of women scientists." The Guardian, July 25. <https://www.theguardian.com/science/occams-corner/2013/jul/25/wikipedia-next-generation-women-scientists>.

<sup>9</sup>Cecco, Leyland. 2018. "Female Nobel prize winner deemed not important enough for Wikipedia entry." The Guardian, October 3. <https://www.theguardian.com/science/2018/oct/03/donna-strickland-nobel-physics-prize-wikipedia-denied>.

increases female students' likelihood of choosing economics as a major. Exposure to female teachers or professors has been shown to increase female students' likelihood of taking math and science courses (Carrell et al., 2010; Lim and Meer, 2019).<sup>10</sup> More generally, Beaman et al. (2009) find that exposure to female leaders weakens stereotypes about gender roles and reduces gender bias. Even though there is no evidence that specifically the lack of detail in Wikipedia about female role models affects economic decisions, more broadly, literature has shown that media affects economic decisions (for an overview see DellaVigna and La Ferrara (2015)) and literature has also shown that availability of information specifically on Wikipedia affects outcomes in the real world (Thompson and Hanley, 2017; Hinnosaar et al., 2017).

In this paper, I study gender inequality in Wikipedia. My goal is to learn why women are less represented and how to change that. Specifically, the questions I ask are the following: (1) Why are women less likely to contribute to Wikipedia? (2) What are the implications of this gender gap on Wikipedia's content? (3) What can be done about it? The analysis also sheds some light on gender inequality in traditional media. It highlights the importance of representing diverse opinions and informs about measures to achieve that.

To answer these questions, I conduct a survey and a randomized survey experiment with 1,000 participants on Amazon Mechanical Turk.<sup>11</sup> This sample is representative of potential editors of Wikipedia. In fact, 22% of the survey respondents have edited Wikipedia. However, there are gender differences—the percentage of Wikipedia editors is almost twice as large among male respondents than female. The survey allows me to study why individuals do or do not contribute to Wikipedia. The goal of the survey was to understand which factors influence Wikipedia editing behavior and how the gender of editors affects editorial input. The goal of the experiment was to test whether providing information about gender inequality in Wikipedia changes editing behavior, specifically, whether the information provided to the treatment group (1) increases the likelihood of editing Wikipedia articles about women and (2) increases the likelihood that survey participants, especially women, edit Wikipedia in the future.

It has been puzzling that women contribute less to Wikipedia than men. There are no large gender differences in the amount of free time.<sup>12</sup> The gender gap in Internet usage is small, and there are no large gender differences in online behaviors such as social networking, online news consumption, communication, and e-commerce.<sup>13</sup> Moreover, women contribute

---

<sup>10</sup>Similarly, racial minority students have been shown to benefit from taking courses with racial minority instructors (Fairlie et al., 2014).

<sup>11</sup>The randomized survey experiment is registered at the AEA RCT Registry with the number AEARCTR-0000500.

<sup>12</sup>Aguiar and Hurst (2007) document that in the US, men and women allocate about equal time to leisure.

<sup>13</sup>Of US adults, 87 percent of men and 86 of women used the Internet in 2014 according to the Pew Research Center ("The Web at 25." 2014, Pew Research Center, February 28. <http://www.pewinternet>).

no less than men to another example of online public good provision, writing user reviews for products and services.<sup>14</sup> Based on the survey, I find that almost half of the gender gap in having edited Wikipedia is explained by gender differences in two characteristics: frequency of Wikipedia use and belief about one’s competence.

The fact that most contributors are men could lead to biased coverage in Wikipedia. I look at the articles about human beings. In the survey experiment, I find that women are almost three times as likely as men to contribute to Wikipedia articles about women. I analyze whether the distribution of editorial input reflects demand for the articles among readers. However, I find no evidence that articles about women have lower demand than those about men.

What can be done about gender inequality in Wikipedia? From a randomized survey experiment, I find that providing information about gender inequality in Wikipedia almost doubles the likelihood of contributing to articles about women. However, the treatment of providing information decreases men’s intention to contribute to Wikipedia and has no effect on women. That is, the treatment leads editors to redirect their editorial input and equalizes intended participation among men and women but decreases overall intended participation in Wikipedia editing. Because there are other ways to increase overall contributions to Wikipedia, the informational treatment can still provide a useful way to decrease gender inequality.

The paper contributes to a number of different areas of research including media bias, user-generated content, and gender differences. Literature in media economics has studied the sources of media and editorial bias.<sup>15</sup> In the case of traditional media, for example, Enikolopov et al. (2011) and Durante and Knight (2012) have provided evidence how certain news outlets give more media coverage to their favored politicians. In this paper, I find that the coverage of public figures in Wikipedia also depends on editors’ characteristics.

---

[org/2014/02/25/the-web-at-25-in-the-u-s](http://org/2014/02/25/the-web-at-25-in-the-u-s)). In the UK in 2014, 82 percent of men and 81 percent of women between 16 and 74 years old used the Internet daily according to Eurostat. In the same sample in the UK, more women participated in online social networks (62 percent of women versus 58 of men), and more men read online news (64 percent of men versus 55 of women). In other activities, such as sending emails, making online purchases, and doing online banking, the differences are even smaller. (“Information Society Statistics,” Eurostat. <http://ec.europa.eu/eurostat/web/information-society/data/database>.)

<sup>14</sup>Chen et al. (2010) found that women contribute more to online movie ratings. Surveys by Pew Research Center found no gender difference in writing user reviews (“Health Online 2013,” 2013, Pew Research Center. [http://www.pewinternet.org/files/old-media//Files/Reports/PIP\\_HealthOnline.pdf](http://www.pewinternet.org/files/old-media//Files/Reports/PIP_HealthOnline.pdf); “Online Product Research,” 2010, Pew Research Center. <http://www.pewinternet.org/files/old-media//Files/Reports/2010/PIP%20Online%20Product%20Research%20final.pdf>).

<sup>15</sup> When analyzing the sources of media bias, the literature has concentrated mainly on traditional media, studying demand-side factors (e.g. Gentzkow and Shapiro (2010); Sen and Yildirim (2015)), the role of advertising (e.g. Di Tella and Franceschelli (2011); Sun and Zhu (2013); Beattie et al. (2017)), and media capture, where government and other special interest groups try to control the content of media (for overviews see Enikolopov and Petrova (2015); Strömberg (2015)).

This is somewhat surprising, considering that it should be more difficult to gain control over Wikipedia than over a single news outlet as anyone is free to edit Wikipedia. Bias in Wikipedia has been studied in the political articles by Greenstein and Zhu (2012, 2014). The current paper is the first to comprehensively analyze the sources of gender inequality in coverage in Wikipedia.

There is an emerging small branch of literature in economics that studies contributions to new media and user-generated content (for an overview, see Luca (2015)). The current paper is most closely related to the studies on contributions to Wikipedia. Most notably, Zhang and Zhu (2011) studied the relationship between contributions to Wikipedia and group size. Aaltonen and Seiler (2015) studied the impact of the cumulative growth of content on new contributions. Algan et al. (2013) conducted online experiments with current Wikipedia editors to test whether pro-social motives explain their contributions to Wikipedia, and Gallus (2017), using a field experiment, estimated the impact of symbolic rewards on Wikipedia editor retention. None of these papers examined the role of gender in contributions, which is the subject of this paper.<sup>16</sup>

Finally, the paper is also related to the literature on gender differences in self-assessment, contributions, and competition.<sup>17</sup> Gender differences in self-assessment have been related to the gender differences in various outcomes, including that women are less likely to give advice in strategic settings (Cooper and Kagel, 2016) and contribute their ideas in group-work settings (Coffman, 2014). The results presented in this paper are consistent with the findings of Coffman (2014), who shows that a big part of the gender differences in contributing ideas is driven by the gender differences in beliefs about competence. Editing Wikipedia resembles a competition where losers' contributions are deleted by fellow editors, while winners' contributions remain. The literature on competition has shown that women are less likely to choose to compete (Niederle and Vesterlund, 2007) and work in competitive settings (Flory et al., 2015).<sup>18</sup> The current paper provides evidence from another competitive setting where women are less likely to choose to participate.

---

<sup>16</sup>In the case of open-source software and related Q&A communities, studies have shown that majority of contributors are men (David and Shapiro, 2008; Vasilescu et al., 2014), women's contributions are valued less (Bohren et al., 2018; Terrell et al., 2017), and women self-promote their skills less (Murciano-Goroff, 2018).

<sup>17</sup>For overviews see Croson and Gneezy (2009); Niederle and Vesterlund (2011); Niederle (2017). Among other things the literature has found that men are more (over-)confident than women and are more likely to over-estimate their probability of success at a given task (see, for example, Beyer (1990); Beyer and Bowden (1997)). The gender differences in self-assessment have been associated with the tendency of women to attribute success more externally, taking less credit for their performance (see, for example, Feather (1969); Deaux and Farris (1977); Beyer (1998)).

<sup>18</sup>The literature has also shown that women perform worse than men in competitive environments even when they perform similarly in non-competitive environments (Gneezy et al., 2003). The gender difference in the willingness to compete is reversed in matrilineal societies (Gneezy et al., 2009) and no gender difference is found when competing against one's own previous performance (Apicella et al., 2017).

The remainder of the paper is organized as follows. Section 2 describes Wikipedia’s editing process and presents basic empirical facts about gender inequality in Wikipedia. Section 3 describes the survey and the randomized survey experiment. Section 4 presents empirical results that answer three questions. First, why are women less likely to edit Wikipedia compared to men? Second, how does the gender gap in participation affect biographies on Wikipedia? Third, how does the treatment of providing information affect gender inequality in Wikipedia? Section 5 concludes.

## 2 Background and basic empirical facts

### 2.1 Wikipedia’s editing process

Wikipedia is a free-access Internet encyclopedia. As of November 2017, it contains about 47 million articles, including 5.5 million articles in English-language Wikipedia.<sup>19</sup> Wikipedia is owned by the Wikimedia Foundation, a non-profit foundation. It is written by volunteers.

Anyone can create Wikipedia articles and edit almost any of its existing articles. There are two types of editors: anonymous and registered. Registration is not required for most editing activities. However, to create new articles or to upload images, editors have to register. For each article, Wikipedia records who edited the article, when, and how. In the case of anonymous editors, their computers’ IP addresses are recorded.

Wikipedia’s editing community is large. According to Wikimedia Statistics,<sup>20</sup> as of September 2017, about 6.7 million registered editors had contributed to the English-language edition of Wikipedia. Most editors made only a few edits, but about 1.2 million editors had edited at least 10 times.

Most editors don’t receive monetary rewards or derive career-related benefits from contributing to Wikipedia (Algan et al., 2013; Gallus, 2017).<sup>21</sup> This is in contrast to open-source software, where developers could be motivated by career concerns (Lerner and Tirole, 2002). Most Wikipedia editors are anonymous or use pseudonyms and each article is a work product of many editors, which makes it difficult for the editors to use their experience on Wikipedia to directly enhance their career prospects. Typically, there is also no personal direct benefit from having better quality Wikipedia articles. Again, this is in contrast to open-source software developers benefiting from improving their own software (Hertel et al., 2003).

---

<sup>19</sup>[https://meta.wikimedia.org/wiki/List\\_of\\_Wikipedias](https://meta.wikimedia.org/wiki/List_of_Wikipedias), accessed November 15, 2017.

<sup>20</sup> <http://stats.wikimedia.org/EN/TablesWikipediaEN.htm>, accessed November 15, 2017.

<sup>21</sup>The exception is paid editing, which is generally discouraged by Wikipedia; since 2014 all contributors are required by the Wikipedia terms of use to disclose any paid editing. Source: [https://en.wikipedia.org/wiki/Conflict-of-interest\\_editing\\_on\\_Wikipedia](https://en.wikipedia.org/wiki/Conflict-of-interest_editing_on_Wikipedia).

## 2.2 Basic empirical facts about gender inequality in Wikipedia

In this section, I describe basic empirical facts about gender inequality in Wikipedia, which provide a broad motivation for the following analysis.

There exists evidence that the majority of Wikipedia editors are men. The main sources about the characteristics of Wikipedia editors are the surveys conducted by or in close collaboration with the Wikimedia Foundation. For example, in the 2018 Wikipedia Community Engagement Insights survey, 85% of English Wikipedia editors who reported their gender identified as male.<sup>22</sup> In earlier surveys, also more than 80 percent of the editors were men.<sup>23</sup> Similar evidence has been found in other studies involving Wikipedia editors. For example, in an experimental study conducted among Wikipedia editors, 90% of the sample were men (Algan et al., 2013). In a survey conducted among top editors of Wikipedia medical articles, 82% of the respondents were men (Heilman and West, 2015). However, less is known about the differences in the male and female editors' editing behavior.

Here, I report basic empirical facts about the editing behavior of registered editors of English language Wikipedia who have reported their gender when registering their Wikipedia username. The caveat of the dataset is that this is a selected sample of editors who voluntarily report their gender. We might be worried that especially women, being afraid of hostility or discrimination, are less likely to publicly reveal their gender. On the other hand, the dataset provides the most detailed information about editors' behavior. I describe their editing behavior of articles about human beings. Concentrating on biographies provides a clear way to study the differences in the editorial input, by analyzing which editors edit articles about men or women. Details of the data are provided in appendix A.1.

Figure 1 presents the average percentage of contributions to biographies of women vs men by the gender of Wikipedia editor. The contributions are measured in the number of biographies edited. Figure 1 shows that female editors (compared to male editors) direct relatively more (31% compared to 17%) of their editorial input towards biographies of women. The difference is significant according to t-test and non-parametric Wilcoxon rank-sum test, both with p-values less than 0.001. Additional analysis in appendix A.2 shows similar results with alternative outcome measures and alternative samples.

The above shows that in this sample, editors tend to contribute to different articles. This

---

<sup>22</sup>Community Engagement Insights 2018 Report, [https://meta.wikimedia.org/wiki/Community\\_Engagement\\_Insights/2018\\_Report/Contributors](https://meta.wikimedia.org/wiki/Community_Engagement_Insights/2018_Report/Contributors).

<sup>23</sup> Wikimedia Foundation. 2011. "Editor Survey," April. [http://upload.wikimedia.org/wikipedia/commons/7/76/Editor\\_Survey\\_Report\\_-\\_April\\_2011.pdf](http://upload.wikimedia.org/wikipedia/commons/7/76/Editor_Survey_Report_-_April_2011.pdf); Wikimedia Foundation. 2011. "Second 2011 Wikipedia Editor Survey," December. [https://upload.wikimedia.org/wikipedia/commons/8/84/December\\_2011\\_Wikipedia\\_Editor\\_Survey\\_topline.pdf](https://upload.wikimedia.org/wikipedia/commons/8/84/December_2011_Wikipedia_Editor_Survey_topline.pdf); Glott, Ruediger; Schmidt, Philipp; Ghosh, Rishab. 2010. "Wikipedia Survey: Overview Results", March. [http://web.archive.org/web/20100414165445/http://wikipediasurvey.org/docs/Wikipedia\\_Overview\\_15March2010-FINAL.pdf](http://web.archive.org/web/20100414165445/http://wikipediasurvey.org/docs/Wikipedia_Overview_15March2010-FINAL.pdf).

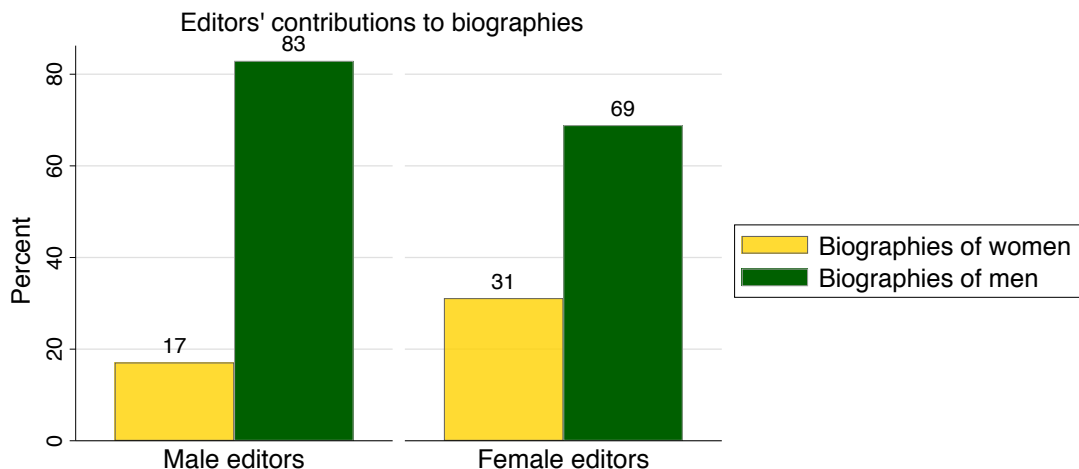


Figure 1: Editors’ contributions to biographies of women vs men, by gender of the editor

Note: The average (over editors) percentage of contributions to the biographies of women vs men, by the gender of Wikipedia editor. Contributions are measured in the number of biographies edited. Calculated based on registered editors who have reported their gender and edited biographies. There are 65,398 such male and 9,419 such female editors. The difference is significant according to t-test and non-parametric Wilcoxon rank-sum test, both with the p-value less than 0.001.

raises the question whose contributions are more likely to correspond to readers’ demand. Is the allocation or editorial input proportional to the readership? To answer this, let’s look at the Wikipedia viewership (see appendix A.3 for further details). Each day, there is a large share of Wikipedia articles that noone reads. On a typical day, the percentage of biographies of men that noone reads is 26% which is larger than that of women (only 16%). Over a longer time period, almost all articles have readers.

Figure 2 describes the distribution of biographies by the average number of views per day (for details of the calculations see appendix A.3). It shows that biographies of men receive fewer views than those of women. For example, 23% of biographies of men and only 13% of biographies of women have on average less than one view per day. On the other end of the distribution, only 2% of biographies of men and 5% of those of women receive 100 or more views per day. Appendix A.3 provides further evidence showing that biographies of women receive relatively more views compared to those of men.

### 3 Survey and experiment

The survey together with the randomized survey experiment was administered in September 2014. The survey had the following structure: (1) background socioeconomic questions



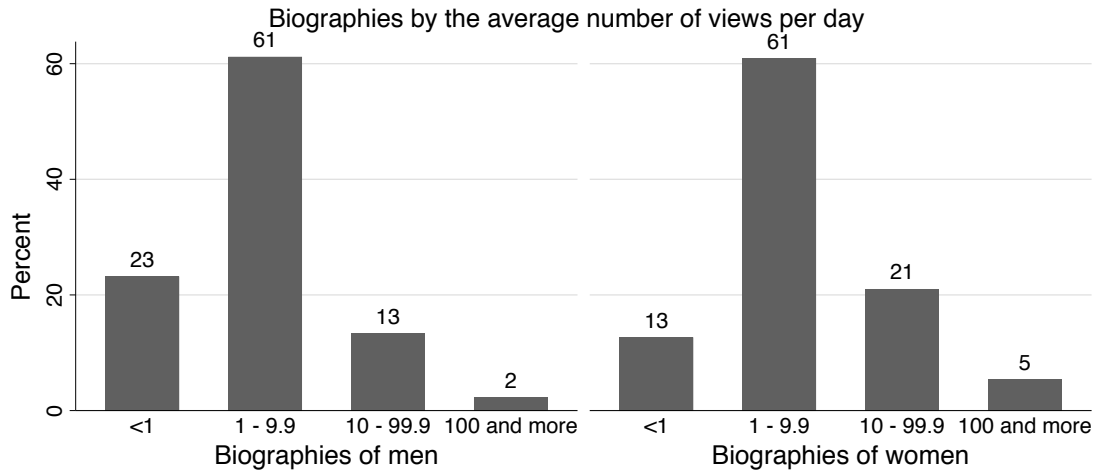


Figure 2: Biographies by the average number of views per day

Note: Calculated based on page views in September 2014 of 784,178 biographies of men and 140,193 biographies of women in the English language Wikipedia. Details of the calculations are in appendix A.3.

including gender, age, education, work status, and children; (2) questions about previous Wikipedia use and editing; (3) randomized treatment providing information about gender inequality in Wikipedia shown only to the treatment group; the control group was shown information about Wikipedia that was not related to gender; (4) hypothetical situation where one is asked to edit Wikipedia; specifically, respondents were asked to choose a Wikipedia page of a person and find some information from the Internet that is missing from the Wikipedia page; note that the respondents were not asked to actually edit Wikipedia; (5) questions about Wikipedia editing, including respondents' beliefs about their competence and the likelihood of editing Wikipedia in the future; and (6) questions about charitable giving, volunteering, writing user reviews, and free time.

In conducting the randomized survey experiment, I follow the approach taken by Kuziemko et al. (2015), who studied how information about inequality changes reported preferences for redistribution in a similar research design on Amazon Mechanical Turk. Amazon Mechanical Turk has recently gained popularity as a platform for experiments in economics (for other examples, see Horton et al., 2011). Methodological papers have evaluated the platform's suitability for surveys and experiments (see, for example, Mason and Suri, 2012), and I followed their advice.

### 3.1 Randomized survey experiment

The goal of the experiment was to test whether providing information about gender inequality in Wikipedia changes editing behavior. Individuals were randomly assigned to either a treatment or a control group with equal probabilities. The information provided to the treatment group is a quote from the Wikipedia page, “Gender Bias on Wikipedia”<sup>24</sup>: “Wikipedia has been criticized by some academics and journalists for having only 9% to 13% female contributors and for having fewer and less extensive articles about women or topics important to women.” The control group read the following information about Wikipedia, which has nothing to do with gender bias: “Wikipedia started in 2001. English-language Wikipedia has over 4.5 million articles.”

After reading the information, respondents were asked to imagine a hypothetical situation in which they edit a person’s Wikipedia page. Respondents were asked to look at Wikipedia articles and find some relevant information from the web that is missing from a Wikipedia article. Note that respondents were never asked to actually edit Wikipedia. Respondents were only asked to report the Wikipedia page they would choose and the information they would hypothetically add to the page. Specifically, the respondents were asked the following questions: (1) “Which human being would you choose? Please write here the name of the human being.” (2) “Please provide a link to the Wikipedia page of the human being, the page that you would choose to edit. Please make sure that it is a valid link to the page of the human being on English language Wikipedia; it must start with: <http://en.wikipedia.org/wiki/>” (3) “Please write here the text that you would add to the page.” (For the screenshot of the instructions and questions, see figure B.1 in online appendix B.) In the end, they were also asked how likely they are to edit Wikipedia in the future.

**Limitations.** Survey experiments are commonly used in various fields in economics to study the impact on stated preferences, willingness to pay, and intentions. But the method raises a concern that the actual behavior might be different from reported intentions. To alleviate the concern, I design the survey experiment to resemble the Wikipedia editing process. Answering the survey question about which page you would edit requires much more effort than just picking a name. A respondent has to go to Wikipedia, choose an article, search on the web for relevant information missing from the article, and report that in the survey. The instructions only stopped short of asking to actually add the information to the Wikipedia article, and this should be the least time-consuming part. Therefore, the data on the reported choices of pages which to edit can be considered reasonably reliable. However,

---

<sup>24</sup> The Wikipedia page, “Gender Bias on Wikipedia”, was from August 24, 2014: [http://en.wikipedia.org/w/index.php?title=Gender\\_bias\\_on\\_Wikipedia&oldid=622670577](http://en.wikipedia.org/w/index.php?title=Gender_bias_on_Wikipedia&oldid=622670577)

the data on the survey respondents' intentions to edit Wikipedia in the future is only about their reported intentions and might not correspond to their actual behavior.

## 3.2 Data collection

The survey was posted on Amazon Mechanical Turk. Amazon Mechanical Turk is an Internet marketplace, where workers perform tasks, including answering surveys and participating in experiments. The survey was described as a 20-minute research study on Wikipedia, which pays the participant \$1.50. The payment gives an effective hourly wage close to the average effective hourly wage on Amazon Mechanical Turk according to Kuziemko et al. (2015).

To ensure high-quality responses, several measures were taken based on advice from previous studies that used Amazon Mechanical Turk. First, only workers with a good track record were allowed to take the survey. The exact qualification criteria were that a worker must have completed and approved at least 50 tasks, and his or her approval rate must be 95 percent or higher. Second, only workers from the US were eligible to take the survey. I chose to limit the set of possible respondents to those in somewhat similar environments because the survey asked about volunteering and writing user reviews. The location was checked and restricted by Amazon itself. According to computer IP addresses of the workers who completed the survey, a very small percentage came from locations outside of the US. Third, respondents were told that payment would be contingent on completing the survey and providing a survey code visible only after finishing. The respondents were not allowed to skip any questions. Fourth, respondents were told that they can start the survey only once. The main concern was that they reach the treatment page multiple times and see both the treatment group's and the control group's information. One respondent reached the treatment page twice, but by chance, he received the same treatment both times, so I still include his responses in the sample.

The survey stayed open until 1,000 workers reported that they had completed it. Survey attrition is analyzed in appendix C. Of the 1,000 Amazon Mechanical Turk workers, 28 did not provide a valid survey response: four workers could not be matched to any survey response (because they either had not started the survey or did not provide enough information to match to the survey response), three started but did not complete the survey, and 21 did not provide a Wikipedia biography of a human being (because they either misunderstood the question or ignored the instructions, for example, chose a fictional character instead of a human being). After excluding these respondents, 972 remained.

**Quality of the data about the choices of pages which to edit.** The main response variables about Wikipedia editing are hypothetical, which raises a concern about the reliability

of the data. As discussed in subsection 3.1, the questions that guide through the Wikipedia editing process are designed to be detailed enough, requiring some effort, to make the data about the page choices reasonably reliable. In the following, let's assess the quality of the answers.

In the hypothetical editing situation, respondents were essentially asked to do two things: (1) to look up and provide a Wikipedia biography of their choice, and (2) to provide some relevant information that is missing from the page. It is clearcut how to evaluate whether an answer was a biography in Wikipedia. Indeed, by the construction of the main sample, all of the 972 respondents provided a biography of a human being.

The evaluation of the quality of the second part of the answer is not straightforward. To get some idea of whether the respondents followed instructions and provided information, the responses are categorized to added information or not. Note that this classification is subjective. Let me illustrate it. Here are some examples of answers that were categorized as "not adding information": about Albert Einstein, "his hair was unconventionally eccentric"; Grace Hopper "is an inspiration to male and female computer scientists alike"; about J. R. R. Tolkien, "one of the best linguists of the 20th century". Answers that were too general or ambiguous were also categorized as "not adding information", for example, about Sojourner Truth, "something more about her family"; about Persius, "I would shorten the scope of what is accredited to him"; about Nicola Tesla, "I would expand the section about his role in the development of electricity". Most of the answers were classified as "adding information". These provided some specific facts, often about personal life, family, and relatives, or details about their work and careers. For example, Abraham Lincoln "created the Secret Service hours before his assassination"; Albert Einstein "had a pet parrot named Bibo". As said above, the classification is subjective and some of the answers categorized as "not providing information" could nevertheless improve Wikipedia.

According to this classification, about 85 percent of the responses provided some information. While the percentage of good responses is rather high, it is subject to the caveat that the decision criteria to evaluate the quality is subjective. Therefore, all the following analysis in the main part of the paper will use the sample of all the 972 respondents that provided a biography in Wikipedia. However, appendix D repeats the analysis restricting the sample to these 85 percent of the responses categorized as providing information and the results are very similar to the ones from the main sample.

### 3.3 Gender and other characteristics

Table 1 presents characteristics of the 972 respondents who completed the survey and compares their characteristics to those from other surveys about Internet users, Wikipedia users, and Wikipedia editors. Column 1 presents demographic characteristics of the survey participants: 49.5 percent of the sample is female, 53.9 percent has at least a college degree, 57.7 percent is employed full-time (works at least 35 hours per week), and 9.9 percent is unemployed. Although the population on Amazon Mechanical Turk excludes individuals under 18-years-old, the sample is rather young, only 35.5 percent is at least 35 years old. Altogether, the sample is similar to other studies using Amazon Mechanical Turk, including Kuziemko et al. (2015).

Table 1: Survey respondents’ characteristics and comparison with other surveys

	Current survey		Comparison surveys		
	All respondents	Only editors	Internet users Pew 2010	Wikipedia users Pew 2010	Wikipedia editors WF 2011
	(1)	(2)	(3)	(4)	(5)
Female	0.495	0.315	0.514	0.485	0.135
At least 35 years old	0.355	0.286	0.628	0.536	0.452
College degree	0.539	0.596	0.344	0.445	0.700
Employed full-time	0.577	0.592	0.513	0.539	0.452
Unemployed	0.099	0.085			
Use Wikipedia daily	0.277	0.460			
Use Wikipedia at least weekly	0.811	0.906			
Has edited Wikipedia	0.219	1.000			
Observations	972	213	1756	437	984

Note: Each cell presents the mean value of a binary characteristic. Columns 1 and 2 present characteristics of the survey in the paper. In column 1, the sample includes all valid survey responses. In column 2, the sample is restricted to those who have edited Wikipedia in the past. For comparison, columns 3–5 present results from other surveys: in columns 3–4, Pew Research Center’s Internet and American Life Project 2010 survey, where Pew survey provided weights are used to correct for the response bias; in column 5, Wikimedia Foundation’s Wikipedia Editors Survey 2011, where the sample is restricted to editors from the U.S.

Column 1 of table 1 also shows that 21.9 percent of the survey sample has edited Wikipedia previously. How close is this percentage to the share of Wikipedia editors in a random sample of Internet users? It might be not too far off because a back-of-the-envelope calculation suggests that in 2010 17 percent of Internet users had edited Wikipedia. The back-of-the-envelope calculation is subject to caveats, because it combines information from two different

surveys, according to which in 2010: 53 percent of U.S. Internet users used Wikipedia<sup>25</sup> and 33 percent of Wikipedia users had edited Wikipedia.<sup>26</sup>

Column 2 of table 1 restricts the sample to those who have edited Wikipedia previously. Of those who have edited, only 31.5 percent are women. Recall that the original sample of survey respondents was gender-balanced (see column 1), but the editors' sample is not. Hence, the survey also provides evidence of gender imbalance among Wikipedia editors.

To understand how representative the current survey sample is of potential and actual Wikipedia editors, table 1 also presents demographic characteristics of Internet users (column 3), Wikipedia users (column 4), and Wikipedia editors (column 5). The data on Internet users and Wikipedia users is from Pew Research Center's Internet and American Life Project, May 2010 survey<sup>27</sup> and the data on Wikipedia editors is from the Wikimedia Foundation's Wikipedia Editors Survey 2011,<sup>28</sup> where the sample is restricted to editors from the U.S. While anyone can edit Wikipedia, Internet users and Wikipedia users are the most likely potential editors. The sample in this survey (column 1) is similar to the Internet and Wikipedia users in terms of gender and employment while being younger and more educated. Those who have edited Wikipedia in the current survey (column 2) compared to Wikipedia editors in the Wikimedia Foundation's survey (column 5) are somewhat younger and more likely to be female, somewhat less likely to have a college degree and less likely to be full-time employed.

Table 2 presents characteristics of men and women in the control group separately by gender. The treatment group is excluded here to make sure the responses are not affected by the treatment, because some of the questions, specifically those that were not included in table 1, were asked after treatment. The table also reports p-values from the t-test for whether the difference between men and women is significantly different from zero.<sup>29</sup> In terms of demographic characteristics, there are significant differences between men and women only in employment: 65 percent of men and 46 percent of women are employed full-time.

Table 2 shows that there are several significant differences between men and women in terms of Wikipedia-related behaviors and attitudes. First, men are almost twice as likely to have edited Wikipedia. The difference is statistically significant with the p-value of the t-test

---

<sup>25</sup>Zickuhr, K. and Rainie, L. 2011. "Wikipedia, past and present", Pew Research Center. <http://www.pewinternet.org/2011/01/13/wikipedia-past-and-present/>, accessed January 23, 2019.

<sup>26</sup>Glott, Ruediger; Schmidt, Philipp; Ghosh, Rishab. 2010. "Wikipedia Survey: Overview Results", March. [http://web.archive.org/web/20100414165445/http://wikipediasurvey.org/docs/Wikipedia\\_Overview\\_15March2010-FINAL.pdf](http://web.archive.org/web/20100414165445/http://wikipediasurvey.org/docs/Wikipedia_Overview_15March2010-FINAL.pdf).

<sup>27</sup>Pew Research Center's Internet and American Life Project, May 2010 survey dataset is available from: <http://www.pewinternet.org/dataset/may-2010-cell-phones/>, accessed Jan 22, 2019.

<sup>28</sup>Wikimedia Foundation's Wikipedia Editors Survey 2011 dataset is available from: <https://dumps.wikimedia.org/other/surveys/editorsurvey2011/>, accessed Jan 22, 2019.

<sup>29</sup>Note that as all the reported characteristics are binary variables, the t-test and the corresponding non-parametric Wilcoxon rank-sum test give the same results.

Table 2: Summary statistics of survey respondents in the control group, split by gender

	Men	Women	p-value
Demographic characteristics			
At least 35 years old	0.320	0.382	0.154
College degree	0.514	0.536	0.611
Works at least 35 hours per week	0.653	0.459	0.000
Unemployed	0.089	0.116	0.322
Wikipedia			
Use Wikipedia daily	0.359	0.176	0.000
Use Wikipedia at least weekly	0.838	0.742	0.009
Has edited Wikipedia	0.313	0.163	0.000
Likely to edit Wikipedia	0.147	0.094	0.077
Believe: more competent than other editors	0.085	0.043	0.059
Believe: at least as competent as other editors	0.761	0.670	0.025
Other			
Donated to charity	0.475	0.554	0.081
Volunteering at least one hour per week	0.332	0.365	0.447
Wrote user review in previous six months	0.517	0.545	0.540
Leisure at least 3 hours per weekday	0.741	0.734	0.852
Observations	259	233	

Note: The sample is restricted to the control group. Columns 1 and 2 present the mean values of the survey respondents’ binary characteristics separately for men and women. Column 3 presents p-values of the t-test for whether the difference between men and women is significantly different from zero. *Likely to edit Wikipedia* is an indicator variable that takes value one if answered “very likely” or “quite likely” to the question “How likely or unlikely is it that you will actually edit Wikipedia over the next 30 days?” *Believe: more competent than other editors* is an indicator variable that takes value 1 if answered “more knowledgeable/competent” to the question “Think of the Wikipedia article of the human being that you chose. On the whole, do you think that you are more or less knowledgeable and competent to edit the article than other people who will edit it in the future?” *Believe: at least as competent as other editors* is an indicator variable that takes value 1 if answered either “more knowledgeable/competent”, “somewhat more”, or “same” to the same question above.

less than 0.001. The difference is not as large as reported in other surveys (see subsection 2.2), but it is still sizable. Second, men use Wikipedia more often — they are twice as likely than women to use Wikipedia daily.<sup>30</sup> Third, a larger share of women believe they are less competent than other editors to edit Wikipedia. Fourth, a larger share of men intend to edit Wikipedia in the future, however the gender difference is rather small, only five percentage points. Note that there are no significant differences between men and women in several other activities like in writing user reviews, in volunteering, and in the amount of free time.

<sup>30</sup> Note that the frequency of Wikipedia use could depend on both the frequency of using encyclopedias and preference for Wikipedia over other sources, such as *Encyclopedia Britannica*.

However, women are significantly more likely than men to give to charity.

Note that this survey is unlikely to be subject to the caveats of women being less likely to respond or less likely to reveal their gender. First, the participation in the survey is gender-balanced. Hence, there is no bias generated by possible non-response. Second, unlike the Wikipedia editing data (subsection 2.2), where editors publicly report their gender and could be afraid of hostility and discrimination, in this survey the participants are anonymous and there is no reason for women to underreport Wikipedia editing. Moreover, in the responses, while there are gender differences in having edited Wikipedia, there are no gender differences in several other behaviors.

## 4 Empirical results

This section provides empirical results that answer three questions. First, why are women less likely to contribute to Wikipedia? Second, what are the implications of the gender gap in participation on the information available on Wikipedia? Third, how does the treatment of providing information about gender inequality in Wikipedia affect editing behavior?

### 4.1 Why are women less likely to edit Wikipedia?

To answer the question of why women contribute less to Wikipedia, I use the survey data restricting the sample to the control group in the survey experiment. I present two sets of results. First, I analyze which characteristics and behaviors are correlated with editing Wikipedia. Second, I analyze how much of the gender difference in Wikipedia editing is explained by the gender differences in these characteristics. I use two measures for Wikipedia editing: (1) whether the respondent has edited Wikipedia in the past and (2) whether the respondent reports that he or she is likely to edit Wikipedia in the future.

Table 3 presents results from logit regressions where the dependent variable, an indicator for editing Wikipedia, is regressed on gender and other characteristics. These include demographic characteristics (indicators for being 35 or older, having a college degree, being unemployed, and working at least 35 hours per week), behaviors and attitudes toward Wikipedia (indicators for using Wikipedia daily, using Wikipedia at least weekly, and belief about one's competence compared to other Wikipedia editors), and other behaviors (indicators for donating to charity in the past year, volunteering at least one hour per week in the past year, writing a user review in the past six months, and having at least three hours of leisure time per weekday).

In columns 1–2, the dependent variable is an indicator for whether the respondent has



Table 3: Estimated role of survey respondents' characteristics in Wikipedia editing

	Has edited Wikipedia		Likely to edit	
	Coef. estimates (1)	Marginal effects (2)	Coef. estimates (3)	Marginal effects (4)
Female	-0.669*** (0.245)	-0.106*** (0.038)	-0.148 (0.317)	-0.014 (0.030)
At least 35 years old	-0.104 (0.251)	-0.017 (0.040)	-0.184 (0.322)	-0.017 (0.030)
College degree	0.429* (0.238)	0.068* (0.037)	0.310 (0.310)	0.029 (0.029)
Unemployed	-0.055 (0.411)	-0.009 (0.065)	0.105 (0.597)	0.010 (0.056)
Works at least 35 hours per week	-0.345 (0.256)	-0.054 (0.040)	0.192 (0.338)	0.018 (0.032)
Use Wikipedia daily	1.025*** (0.254)	0.162*** (0.038)	1.156*** (0.331)	0.109*** (0.031)
Use Wikipedia at least weekly	0.475 (0.358)	0.075 (0.056)	0.063 (0.468)	0.006 (0.044)
Believe: more competent than other editors	0.782* (0.415)	0.124* (0.065)	0.414 (0.479)	0.039 (0.045)
Believe: at least as competent as other editors	0.743** (0.297)	0.117** (0.046)	1.946*** (0.614)	0.184*** (0.059)
Donated to charity	0.238 (0.253)	0.038 (0.040)	0.387 (0.336)	0.037 (0.032)
Volunteering at least one hour per week	-0.146 (0.256)	-0.023 (0.040)	0.345 (0.322)	0.033 (0.030)
Wrote user review in previous six months	0.552** (0.236)	0.087** (0.037)	0.294 (0.306)	0.028 (0.029)
Leisure at least 3 hours per weekday	-0.025 (0.260)	-0.004 (0.041)	0.094 (0.338)	0.009 (0.032)
Log-likelihood	-238.921		-156.098	
Observations	492		492	

Note: Columns 1 and 3 present coefficient estimates from logit regressions and columns 2 and 4 the corresponding marginal effects from the same regressions. In columns 1-2, the dependent variable is an indicator for whether the individual has edited Wikipedia in the past. In columns 3-4, the dependent variable is whether the individual reports he or she is likely to edit Wikipedia in the next 30 days, specifically, the variable takes value one if the respondent answered that he or she is either “very likely” or “quite likely” to the question: “How likely or unlikely is it that you will actually edit Wikipedia over the next 30 days?” Explanatory variables are described by table 2. The sample is restricted to the control group. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

edited Wikipedia in the past. Column 1 presents coefficient estimates and column 2 the corresponding marginal effects. Estimates show that those who use Wikipedia daily are 16 percentage points more likely to have edited Wikipedia. Those who have written a user review in the past six months are 9 percentage points more likely to have edited Wikipedia. Those with a stronger belief about one’s competence compared to other potential editors are also more likely to have edited. Note that the gender difference in the belief about competence could be due to women being less competent or due to women underestimating their competence. The survey data does not allow these to be distinguished.

In columns 3–4, the dependent variable is an indicator for whether the respondent reports he or she is likely to edit Wikipedia in the next 30 days. Specifically, the variable takes value one if the respondent answered that he or she is either “very likely” or “quite likely” to the question: “How likely or unlikely is it that you will actually edit Wikipedia over the next 30 days?” The only characteristics that are statistically significantly associated with the intention to edit in the future are daily use of Wikipedia and belief about one’s competence. Appendix B presents table B.2 with the same regressions estimated separately for men and women.

Let’s address the role of the editors’ gender in past and intended editing. The estimated marginal effects in column 2 in table 3 imply that women are 11 percentage points less likely to have edited Wikipedia in the past even controlling for other factors. But according to the estimated marginal effects in column 4, gender doesn’t significantly affect the intention to edit in the future. This implies that the observed gender gap in the intention to edit (see table 2) is explained by other characteristics. Note that the two editing variables differ not only in terms of measuring past behavior vs future intent but also in terms of frequency of editing. Specifically, the past editing measure is cumulative (“have you ever edited”), while the future intent measures intent to edit in the next 30 days. As most editors make less than ten edits, a person who edits once a month is already a frequent Wikipedia editor. That gender plays a different role in the case of these two editing measures is consistent with the explanation that the two measures capture different parts of the distribution of editors (more or less frequent editors) and that only in the case of less frequent editing, gender conditional on other factors plays a significant role.

To learn how much of the gender difference in Wikipedia editing (both past and intended) is explained by gender and how much comes from differences in attitudes and behaviors toward Wikipedia, I decompose the gender difference in editing Wikipedia using Oaxaca

(1973) decomposition, which was extended to non-linear equations by Fairlie (2005).<sup>31</sup> For the explanatory variables, I include indicators for daily Wikipedia use, at least weekly Wikipedia use, and belief about one’s competence compared to other editors.

Table 4 presents the results from the non-linear decomposition. In the first column, the dependent variable is an indicator of whether the respondent has edited Wikipedia in the past. The difference in the fractions of men and women having edited in the past is 15.0 percentage points. Gender differences in the frequency of Wikipedia use and the belief about one’s competence explain a large share, namely 43 percent of the gender gap in past editing. In the second column, the dependent variable is whether the individual is likely to edit Wikipedia in the next 30 days. The gender difference in the intention to edit is 5.2 percentage points. Gender differences in the same characteristics explain 79 percent of the gender gap in the intention to edit Wikipedia in the future.

## 4.2 The effects of the gender gap on Wikipedia content

In this subsection, I study whether the gender gap among editors could lead to gender inequality in the biographical information available on Wikipedia. Recall that in the hypothetical survey experiment, each survey respondent picked a biography to edit. I test whether women are more likely than men to choose to edit biographies of women. In the analysis in this part, the sample is restricted to the control group.

Table 5 presents the share of respondents, separately by gender, who chose to edit a biography of a woman. The table also presents p-values from the t-test of whether the difference between the choices of men and women is significantly different from zero. The first row shows that 31 percent of female respondents chose to edit a biography of a woman, and only 11 percent of male respondents did that. The difference is statistically significant with the p-value from the t-test less than 0.001. Note that 31 percent of female respondents

---

<sup>31</sup> The decomposition for the non-linear equation  $Y = F(X\hat{\beta})$  can be written as

$$\begin{aligned} \bar{Y}^M - \bar{Y}^W &= \left[ \sum_{i=1}^{N^M} \frac{F(X_i^M \hat{\beta}^P)}{N^M} - \sum_{i=1}^{N^W} \frac{F(X_i^W \hat{\beta}^P)}{N^W} \right] \\ &+ \left[ \sum_{i=1}^{N^M} \frac{F(X_i^M \hat{\beta}^M)}{N^M} - \sum_{i=1}^{N^M} \frac{F(X_i^M \hat{\beta}^P)}{N^M} \right] + \left[ \sum_{i=1}^{N^W} \frac{F(X_i^W \hat{\beta}^W)}{N^W} - \sum_{i=1}^{N^W} \frac{F(X_i^W \hat{\beta}^P)}{N^W} \right] \end{aligned}$$

where  $\bar{Y}^j$  is the average probability of the binary outcome variable  $Y$  by gender  $j = M, W$ ,  $N^j$  is the sample size by gender,  $F$  is the cumulative distribution function of the logistic distribution,  $X_i^j$  is a row vector of explanatory variables of individual  $i$  of gender  $j$ ,  $\hat{\beta}^j$  is a vector of coefficient estimates for gender  $j$ , and  $\hat{\beta}^P$  is a vector of coefficient estimates from a pooled sample of both genders. The term in the first brackets captures the part of the gender gap in editing that is explained by the gender differences in explanatory variables.

Table 4: Decomposition of gender gap in Wikipedia editing

Dependent variable:	Has edited Wikipedia	Likely to edit
	(1)	(2)
Total gap	0.1497	0.0523
Explained gap	0.0642	0.0412
Explained percentage	42.9	78.8
Observations	492	492

Note: Each column presents a separate decomposition. In column 1, the dependent variable is an indicator for having edited Wikipedia in the past. In column 2, the dependent variable is an indicator for whether the individual reports he or she is likely to edit Wikipedia in the next 30 days, defined in the same way as in table 2. There are four explanatory variables: an indicator for daily Wikipedia use, an indicator for at least weekly Wikipedia use, *Believe: more competent than other editors*, and *Believe: at least as competent as other editors*, defined in the same way as in table 3. In the rows, *Total gap* is the difference between the average outcome variable of men and women; *Explained gap* is the difference between the outcome variable of men and women that is explained by the differences in these explanatory variables; *Explained percentage* is the percentage of the *Total gap* that is explained by the differences in these explanatory variables. The sample consists of the survey respondents in the control group.

choosing to edit a biography of a woman is not a small percentage, because they had to choose from the existing biographies in Wikipedia and only 15.2 percent of those are biographies of women. Rows 2 - 7 of table 5 present the statistic by survey respondents' education, age, and whether they have edited Wikipedia in the past. In all the groups, female respondents are more likely to edit biographies of women. The last rows of the table restrict attention to subgroups of biographies (as defined in table B.1): first, lawyers, politicians, scientists, etc. (group Professionals); and second, writers, painters, composers, etc. (group Culture). A large gender difference remains in both groups.

Next, I test whether gender has a statistically significant impact on the likelihood of choosing to edit biographies of women versus biographies of men, even when controlling for other characteristics of the respondent. Table 6 presents estimates from a logit regression where the dependent variable is an indicator for whether the respondent chose to edit an article about a woman. The explanatory variables included are an indicator for being at least 35 years old, having a college degree, being unemployed, working at least 35 hours per week, using Wikipedia daily, using Wikipedia at least weekly, and having edited Wikipedia in the past. Column 1 presents coefficient estimates and column 2 the corresponding marginal effects. The results show that women were 21 percentage points more likely than men to edit an article about a woman versus a man; no other included characteristic has a significant impact.

As in subsection 2.2, we could ask whether the allocation of editorial input of the survey

Table 5: Likelihood of choosing to edit a Wikipedia article about a woman vs a man, by the gender of survey respondent

	Chose an article about a woman		t-test	Obs.
	Male respondents	Female respondents	p-value	
All respondents in the control group				
Total	0.112	0.305	0.000	492
Subsets of respondents by demographic characteristics				
Has edited Wikipedia	0.148	0.395	0.003	119
Hasn't edited Wikipedia	0.096	0.287	0.000	373
College	0.128	0.272	0.004	258
No college	0.095	0.343	0.000	234
Less than 35 years of age	0.084	0.270	0.001	172
At least 35 years of age	0.125	0.326	0.000	320
Subsets of respondents by type of Wikipedia articles				
Biographies: Professionals	0.076	0.205	0.015	165
Biographies: Culture	0.123	0.279	0.016	159

Note: A unit of observation is a survey respondent. Each survey respondent chose a Wikipedia article either about a man or a woman. Columns 1 and 2 present the mean of the indicator variable that the chosen article was about a woman, separately for male and female survey respondents. Column 3 presents p-values of the t-test for whether the difference between male and female respondents is significantly different from zero. Column 4 presents the number of observations. In the first row, the sample consists of the survey respondents in the control group. In the following rows, the sample is further restricted by the survey respondents' demographic characteristics that are listed in the leftmost column. In the last two rows, the sample is restricted to the subgroups of biographies: group *Professionals* includes lawyers, politicians, scientists, etc; and group *Culture* includes writers, painters, composers, etc; the exact definition of groups is in table B.1.

respondents corresponds to readers' demand? Appendix E.1 shows that there is no evidence that the biographies of women chosen by the survey respondents have lower viewership than those of men. There is evidence that images drive viewership in Wikipedia (Nagaraj, 2017). Table E.2 in appendix E.1 shows that while pages with more images do indeed get more views, then even controlling for the number of images on the page, biographies of women still get more views.

### 4.3 How to reduce gender inequality in Wikipedia?

There are many potential ways to reduce gender inequality in Wikipedia. I look at only one of these, which is providing information. The main reason to concentrate on this one is that it is cheap and simple to implement.

This section presents evidence from a randomized survey experiment on the impact of providing information about gender inequality in Wikipedia. First, I look at how the

Table 6: Estimated role of survey respondents’ characteristics on the likelihood of choosing to edit a Wikipedia article about a woman vs a man

	Chose an article about a woman	
	Coefficients (1)	Marginal effects (2)
Female	1.368*** (0.261)	0.206*** (0.037)
At least 35 years old	-0.308 (0.253)	-0.046 (0.038)
College degree	-0.109 (0.239)	-0.016 (0.036)
Unemployed	-0.325 (0.415)	-0.049 (0.062)
Works at least 35 hours per week	-0.165 (0.252)	-0.025 (0.038)
Use Wikipedia daily	0.138 (0.292)	0.021 (0.044)
Use Wikipedia at least weekly	0.055 (0.305)	0.008 (0.046)
Has edited Wikipedia	0.460 (0.282)	0.069 (0.042)
Log-likelihood	-230.888	
Observations	492	

Note: Estimation results are from a logit regression. Column 1 presents coefficient estimates and column 2 presents estimated marginal effects. Dependent variable is an indicator for whether the respondent chose to edit an article about a woman. The explanatory variables are the characteristics of the survey respondent. The sample is restricted to the control group. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

treatment affected respondents’ answers regarding their choice of which Wikipedia article to edit. Specifically, I’m interested in the gender of the person whose biography they chose to edit. Second, I look at whether the treatment affects answers about their intention to edit Wikipedia in the future.

### 4.3.1 Choice of pages to edit

Table 7 presents the effect of the treatment on the answers about the choice of the Wikipedia article respondents would edit. Column 1 shows that the treatment is associated with a 94 percent increase in the share of women’s biographies. Note that in the control group, the share of women’s biographies is only 20 percent. With the treatment, it gets close to

40 percent. The inclusion of demographic characteristics (column 2) decreases the scaled treatment effect from 94 to 89 percent.

Table 7: The effect of treatment on the likelihood of choosing to edit a Wikipedia article about a woman

	Total		Male respondents		Female respondents	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.190***	0.180***	0.129***	0.129***	0.232***	0.234***
	(0.029)	(0.028)	(0.034)	(0.034)	(0.044)	(0.044)
Covariates	No	Yes	No	Yes	No	Yes
Control group mean	0.203	0.203	0.112	0.112	0.305	0.305
Scaled treatment effect	0.937	0.885	1.156	1.154	0.760	0.767
Observations	972	972	491	491	481	481

Note: Each column presents estimates from a separate linear probability model. Dependent variable is an indicator for whether the respondent chose to edit an article about a woman. Covariates include demographic characteristics from table 2. The sample includes both the control and treatment group of the survey respondents. The sample is restricted to men in columns 3-4 and women in columns 5-6. *Scaled treatment effect* measures the percentage change of the share of women’s pages. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

The scaled treatment effect is larger for men than women as shown in columns 3-6. Column 4 shows that the treatment in the case of men is associated with a 115 percent increase in the share of women’s biographies. Column 6 shows that the treatment in the case of women is associated with a 77 percent increase in the share of women’s biographies. The larger treatment effect for male respondents reflects the fact that the starting point for men was lower. In the control group, only 11 percent of male respondents would edit biographies of women compared to 31 percent of female respondents.

Overall, the treatment redirects editorial input and almost doubles the likelihood of contributing to an article about a woman. Appendix E.2 describes how the treatment affects the allocation of editorial input in terms of readership. There is no evidence that the treatment makes the allocation significantly worse or better in terms of readership.

### 4.3.2 Participation

Table 8 presents the effect of the treatment on the answers about the intention to edit Wikipedia in the next 30 days. Column 1 shows that the treatment is associated with a 35 percent decrease in the intention to edit Wikipedia. When including demographic characteristics (column 2), the decrease is 37 percent. That the treatment decreases the intention to edit in the future is somewhat unexpected. Columns 3-6 look at the effect

of the treatment on men and women separately, and this helps to clarify the finding. We see that the treatment is associated with a decrease in the intention to edit in the case of men (columns 3-4), and it has no effect on women (columns 5-6). Namely, column 3 shows that the treatment in the case of men is associated with a 62 percent decrease in intended editing. When including demographic characteristics (column 4), the decrease is 64 percent. To analyze the robustness of the finding, I restrict the sample using three characteristics: Wikipedia use, beliefs about competence, and writing user reviews. The results from the restricted samples are presented in table B.3 in appendix B and are similar to those from the main sample.

Table 8: The effect of treatment on the intention to edit Wikipedia in the future

	Total		Male respondents		Female respondents	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	-0.043**	-0.045**	-0.091***	-0.093***	0.006	0.003
	(0.019)	(0.019)	(0.027)	(0.027)	(0.027)	(0.027)
Covariates	No	Yes	No	Yes	No	Yes
Control group mean	0.122	0.122	0.147	0.147	0.094	0.094
Scaled treatment effect	-0.351	-0.367	-0.618	-0.636	0.068	0.028
Observations	972	972	491	491	481	481

Note: Each column presents estimates from a separate linear probability model. Dependent variable is an indicator for whether the individual believes he or she is likely to edit Wikipedia in the next 30 days, defined in the same way as in table 2. Covariates include demographic characteristics from table 2. The sample includes both the control and treatment group of the survey respondents. The sample is restricted to men in columns 3-4 and women in columns 5-6. *Scaled treatment effect* measures the percentage change of the share of respondents who believe it is likely they will edit Wikipedia. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

The information that the majority of Wikipedia editors are men, leads men to reduce their intended editing effort, but it does not change the behavior of women. The informational treatment in this paper can be considered a moral reminder and also a peer information intervention (information about what peers typically do). Note that it's less costly for men to react to the informational treatment than for women. Men can reduce the gender gap among editors by editing less, and this is easy to do in terms of effort; but for women to reduce gender inequality, the solution is more costly, they should edit more. That women don't report an increase in intended editing is in line with the literature that has found that peer information intervention moves behavior closer to the peer group average.

The treatment, emphasizing the importance of equality, was intended to motivate women to contribute. But it also discouraged men from contributing. More generally, the results



highlight a potential problem with a policy that motivates a minority to contribute while discouraging the majority—the majority group is larger, and therefore, the overall effect on the quantity is likely to be negative. The implication is that policies which encourage contributions from specific groups while potentially discouraging others, should be combined with those that increase overall contributions.

## 5 Conclusions

Using survey data and experimental evidence, I study the reasons why women are less likely to contribute to Wikipedia, the implications of this gender gap, and what can be done about it. The findings are important for a widely used reference source like Wikipedia and there are wider implications for both traditional and new media.

I find that men and women tend to contribute to different articles. This implies that the gender gap among contributors leads to an unequal coverage of topics. Moreover, I find that almost half of the gender difference in having edited Wikipedia is explained by gender differences in two characteristics: frequency of Wikipedia use and belief about one’s competence. Combining the difference in Wikipedia use with the finding that men and women contribute to different topics, it points toward possible equilibrium effects. If a media channel covers fewer topics that women are interested in, then women use it less frequently and are less interested in contributing. On the other hand, those who contribute to Wikipedia may over time start to use Wikipedia even more and grow more confident about their ability to contribute.

The results of the paper suggest that providing information to the editors about gender inequality can alleviate the gender gap in contributions. The informational treatment changes the allocation of editorial input. However it comes at the cost of decreasing the editorial input of men. The result provides an example where encouraging gender equality can partially backfire. Wikipedia has set a goal to increase the share of female editors. One way to achieve this is by discouraging male editors. However, this might not be desirable, especially considering the finding by Aaltonen and Seiler (2015) that past contributions tend to motivate further contributions in Wikipedia. The implication for Wikipedia and other forms of media is that it is important to balance the efforts of attracting new contributors and keeping the current ones.

The paper highlights a trade-off between equality and quantity. The treatment generated opposing effects on the benefits from Wikipedia to the readers. On one hand, increasing equality of contributions is likely to be beneficial to the readers—readership statistics (in section 2) suggest that biographies of women have more readers than those of men and

hence, re-directing editorial input is likely to be beneficial. On the other hand, decreasing the quantity of contributions decreases Wikipedia’s benefit to the readers. The effects on the costs of providing Wikipedia are even more complicated. For example, how do the payoffs of Wikipedia editors depend on fellow editors and the readership of articles; do the editors form correct expectations about the readership? Quantifying the costs and benefits provides avenues for future research.

The paper analyzed gender differences in editing biographies. Literature on competition has found that gender differences are eliminated when competing on behalf of others (Cassar et al., 2016). Contributing to Wikipedia to alleviate the gender gap in biographies could be interpreted as competing on behalf of others. If the result on competition extends to Wikipedia editing, we would expect that in editing more neutral articles compared to biographies, women might be even less likely to participate.

This paper has concentrated on gender, but more broadly it highlights the importance of representing diverse opinions. Similar issues arise in the case of other demographic characteristics. For example, it has been noted that information about black history is lacking in Wikipedia and the Wikimedia Foundation is trying to engage new editors in order to change that.<sup>32</sup>

The analysis in this paper faces limitations. First, using the survey experiment, instead of actual choices, raises concerns. Although the survey experiment resembled Wikipedia editing, participants were never asked to actually edit Wikipedia. It is reassuring that the quality of the answers is rather good (as described in subsection 3.2). Second, the survey experiment used only one specific informational treatment. It remains for future research to analyze whether other treatments are more effective in increasing women’s contributions to Wikipedia.

## References

- AALTONEN, A. AND S. SEILER (2015): “Cumulative Growth in User-Generated Content Production: Evidence from Wikipedia,” *Management Science*, 62, 2054–2069.
- AGUIAR, M. AND E. HURST (2007): “Measuring Trends in Leisure: The Allocation of Time Over Five Decades,” *The Quarterly Journal of Economics*, 122, 969–1006.
- ALGAN, Y., Y. BENKLER, M. F. MORELL, AND J. HERGUEUX (2013): “Cooperation in

---

<sup>32</sup>Smith, Jada F. 2015. “Howard University Fills in Wikipedia’s Gaps in Black History.” New York Times, February 19. <http://www.nytimes.com/2015/02/20/us/at-howard-a-historically-black-university-filling-in-wikipedias-gaps-in-color.html>.

- a Peer Production Economy: Experimental Evidence from Wikipedia,” *Working Paper. NBER Workshop on the Economics of IT and Digitization*.
- APICELLA, C. L., E. E. DEMIRAL, AND J. MOLLERSTROM (2017): “No Gender Difference in Willingness to Compete When Competing against Self,” *American Economic Review: Papers and Proceedings*, 107, 136–140.
- BEAMAN, L., R. CHATTOPADHYAY, E. DUFLO, R. PANDE, AND P. TOPALOVA (2009): “Powerful Women: Does Exposure Reduce Bias?” *The Quarterly Journal of Economics*, 124, 1497–1540.
- BEAMAN, L., E. DUFLO, R. PANDE, AND P. TOPALOVA (2012): “Female Leadership Raises Aspirations and Educational Attainment for Girls: A Policy Experiment in India,” *Science*, 335, 582–586.
- BEATTIE, G., R. DURANTE, B. KNIGHT, AND A. SEN (2017): “Advertising Spending and Media Bias: Evidence from News Coverage of Car Safety Recalls,” Working Paper 23940, National Bureau of Economic Research.
- BEYER, S. (1990): “Gender differences in the accuracy of self-evaluations of performance,” *Journal of Personality and Social Psychology*, 59, 960–970.
- (1998): “Gender differences in causal attributions by college students of performance on course examinations,” *Current Psychology*, 17, 346–358.
- BEYER, S. AND E. M. BOWDEN (1997): “Gender Differences in Self-Perceptions: Convergent Evidence from Three Measures of Accuracy and Bias,” *Personality and Social Psychology Bulletin*, 23, 157–172.
- BOHREN, J. A., A. IMAS, AND M. ROSENBERG (2018): “The Dynamics of Discrimination: Theory and Evidence,” *manuscript*.
- CARRELL, S. E., M. E. PAGE, AND J. E. WEST (2010): “Sex and Science: How Professor Gender Perpetuates the Gender Gap,” *The Quarterly Journal of Economics*, 125, 1101–1144.
- CASSAR, A., F. WORDOFA, AND Y. J. ZHANG (2016): “Competing for the benefit of offspring eliminates the gender gap in competitiveness,” *Proceedings of the National Academy of Sciences*, 113, 5201–5205.

- CHEN, Y., F. M. HARPER, J. KONSTAN, AND S. X. LI (2010): “Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens,” *American Economic Review*, 100, 1358–98.
- COFFMAN, K. B. (2014): “Evidence on Self-Stereotyping and the Contribution of Ideas,” *The Quarterly Journal of Economics*, 129, 1625–1660.
- COOPER, D. J. AND J. H. KAGEL (2016): “A failure to communicate: an experimental investigation of the effects of advice on strategic play,” *European Economic Review*, 82, 24–45.
- CROSON, R. AND U. GNEEZY (2009): “Gender Differences in Preferences,” *Journal of Economic Literature*, 47, 448–474.
- DAVID, P. A. AND J. S. SHAPIRO (2008): “Community-based production of open-source software: What do we know about the developers who participate?” *Information Economics and Policy*, 20, 364–398.
- DEAUX, K. AND E. FARRIS (1977): “Attributing causes for one’s own performance: The effects of sex, norms, and outcome,” *Journal of Research in Personality*, 11, 59–72.
- DELLAVIGNA, S. AND E. LA FERRARA (2015): “Economic and Social Impacts of the Media,” in *Handbook of Media Economics*, ed. by S. P. Anderson, J. Waldfogel, and D. Strömberg, North-Holland, vol. 1 of *Handbook of Media Economics*, 723–768.
- DI TELLA, R. AND I. FRANCESCHELLI (2011): “Government Advertising and Media Coverage of Corruption Scandals,” *American Economic Journal: Applied Economics*, 3, 119–151.
- DURANTE, R. AND B. KNIGHT (2012): “Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi’s Italy,” *Journal of the European Economic Association*, 10, 451–481.
- ENIKOLOPOV, R. AND M. PETROVA (2015): “Chapter 17 - Media Capture: Empirical Evidence,” in *Handbook of Media Economics*, ed. by S. P. Anderson, J. Waldfogel, and D. Strömberg, North-Holland, vol. 1 of *Handbook of Media Economics*, 687–700.
- ENIKOLOPOV, R., M. PETROVA, AND E. ZHURAVSKAYA (2011): “Media and Political Persuasion: Evidence from Russia,” *American Economic Review*, 101, 3253–3285.
- FAIRLIE, R. W. (2005): “An extension of the Blinder-Oaxaca decomposition technique to logit and probit models,” *Journal of Economic and Social Measurement*, 30, 305–316.

- FAIRLIE, R. W., F. HOFFMANN, AND P. OREOPOULOS (2014): “A Community College Instructor Like Me: Race and Ethnicity Interactions in the Classroom,” *American Economic Review*, 104, 2567–2591.
- FEATHER, N. T. (1969): “Attribution of responsibility and valence of success and failure in relation to initial confidence and task performance,” *Journal of Personality and Social Psychology*, 13, 129–144.
- FLORY, J. A., A. LEIBBRANDT, AND J. A. LIST (2015): “Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions,” *The Review of Economic Studies*, 82, 122–155.
- GALLUS, J. (2017): “Fostering Public Good Contributions with Symbolic Awards: A Large-Scale Natural Field Experiment at Wikipedia,” *Management Science*, 63, 3999–4015.
- GENTZKOW, M. AND J. M. SHAPIRO (2010): “What Drives Media Slant? Evidence From U.S. Daily Newspapers,” *Econometrica*, 78, 35–71.
- GNEEZY, U., K. L. LEONARD, AND J. A. LIST (2009): “Gender Differences in Competition: Evidence From a Matrilineal and a Patriarchal Society,” *Econometrica*, 77, 1637–1664.
- GNEEZY, U., M. NIEDERLE, AND A. RUSTICHINI (2003): “Performance in Competitive Environments: Gender Differences,” *The Quarterly Journal of Economics*, 118, 1049–1074.
- GREENSTEIN, S. AND F. ZHU (2012): “Is Wikipedia Biased?” *American Economic Review: Papers and Proceedings*, 102, 343–348.
- (2014): “Do Experts or Collective Intelligence Write with More Bias? Evidence from Encyclopedia Britannica and Wikipedia,” *Harvard Business School Working Paper*, 15-023.
- HEILMAN, J. M. AND A. G. WEST (2015): “Wikipedia and Medicine: Quantifying Readership, Editors, and the Significance of Natural Language,” *Journal of Medical Internet Research*, 17, e62.
- HERTEL, G., S. NIEDNER, AND S. HERRMANN (2003): “Motivation of software developers in Open Source projects: an Internet-based survey of contributors to the Linux kernel,” *Research Policy*, 32, 1159–1177.
- HINNOSAAR, M., T. HINNOSAAR, M. KUMMER, AND O. SLIVKO (2017): “Wikipedia Matters,” SSRN Scholarly Paper ID 3046400, Social Science Research Network, Rochester, NY.

- HORTON, J. J., D. G. RAND, AND R. J. ZECKHAUSER (2011): “The online laboratory: conducting experiments in a real labor market,” *Experimental Economics*, 14, 399–425.
- KUZIEMKO, I., M. I. NORTON, E. SAEZ, AND S. STANTCHEVA (2015): “How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments,” *American Economic Review*, 105, 1478–1508.
- LERNER, J. AND J. TIROLE (2002): “Some Simple Economics of Open Source,” *The Journal of Industrial Economics*, 50, 197–234.
- LIM, J. AND J. MEER (2019): “Persistent Effects of Teacher-Student Gender Matches,” *Journal of Human Resources*, 0218–9314R4.
- LUCA, M. (2015): “Chapter 12 - User-Generated Content and Social Media,” in *Handbook of Media Economics*, ed. by S. P. Anderson, J. Waldfogel, and D. Strömberg, North-Holland, vol. 1 of *Handbook of Media Economics*, 563–592.
- MASON, W. AND S. SURI (2012): “Conducting behavioral research on Amazon’s Mechanical Turk,” *Behavior Research Methods*, 44, 1–23.
- MURCIANO-GOROFF, R. (2018): “Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers,” *manuscript*.
- NAGARAJ, A. (2017): “Does Copyright Affect Reuse? Evidence from Google Books and Wikipedia,” *Management Science*.
- NIEDERLE, M. (2017): “A Gender Agenda: A Progress Report on Competitiveness,” *American Economic Review*, 107, 115–119.
- NIEDERLE, M. AND L. VESTERLUND (2007): “Do Women Shy Away From Competition? Do Men Compete Too Much?” *The Quarterly Journal of Economics*, 122, 1067–1101.
- (2011): “Gender and Competition,” *Annual Review of Economics*, 3, 601–630.
- OAXACA, R. (1973): “Male-Female Wage Differentials in Urban Labor Markets,” *International Economic Review*, 14, 693.
- PORTER, C. AND D. SERRA (2017): “Gender differences in the choice of major: The importance of female role models,” *manuscript*.
- REAGLE, J. AND L. RHUE (2011): “Gender Bias in Wikipedia and Britannica,” *International Journal of Communication*, 5, 21.

- SEN, A. AND P. YILDIRIM (2015): “Clicks Bias in Editorial Decisions: How Does Popularity Shape Online News Coverage?” SSRN Scholarly Paper ID 2619440, Social Science Research Network, Rochester, NY.
- STRÖMBERG, D. (2015): “Media and Politics,” *Annual Review of Economics*, 7, 173–205.
- SUN, M. AND F. ZHU (2013): “Ad Revenue and Content Commercialization: Evidence from Blogs,” *Management Science*, 59, 2314–2331.
- TERRELL, J., A. KOFINK, J. MIDDLETON, C. RAINEAR, E. MURPHY-HILL, C. PARNIN, AND J. STALLINGS (2017): “Gender differences and bias in open source: pull request acceptance of women versus men,” *PeerJ Computer Science*, 3, e111.
- THOMPSON, N. AND D. HANLEY (2017): “Science Is Shaped by Wikipedia: Evidence from a Randomized Control Trial,” SSRN Scholarly Paper ID 3039505, Social Science Research Network, Rochester, NY.
- VASILESCU, B., A. CAPILUPPI, AND A. SEREBRENIK (2014): “Gender, Representation and Online Participation: A Quantitative Study,” *Interacting with Computers*, 26, 488–511.
- ZHANG, X. AND F. ZHU (2011): “Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia,” *The American Economic Review*, 101, 1601–1615.

# A Online Appendix: Data overview and analysis supporting subsection 2.2

## A.1 Dataset of Wikipedia editing histories

The dataset is constructed from editing histories of the English-language edition of Wikipedia. To analyze editing by the gender of the editor, I restrict attention to the editors who have reported their gender when registering their Wikipedia username. I analyze their editing behavior of articles about human beings. In particular, I'm interested in whether they edit articles about men or women.

To learn whether a Wikipedia article is about a human being, I use information from Wikidata, which is a database managed by the Wikimedia Foundation. From Wikidata, I also learn the gender and year of birth of the human being. I include in the sample all the human beings in Wikidata as of July 2014 who were born in any year from 1000 to 2000 and who were either male or female. There are almost one million such human beings. When there was more than one birth year in the database, the earliest was taken (there were 1,445 such human beings in the database). I exclude all human beings whose gender is not determined to be either male or female in Wikidata (either gender is missing, it is not male or female, or the person has more than one gender in the database), there are 41 such cases. After excluding those, we are left with the Wikidata sample of 924,371 human beings.

From the Wikidata database, I also obtain information about the professions of the human beings. Of the sample of human beings selected above, 490,023 have at least one profession in the Wikidata database. I categorize all the professions that are held by at least 1,000 human beings in the sample, which is 88% of all the professions held by these people, into three groups. The first group I call *Professionals*, this includes lawyers, politicians, and scientists among others. The second group I call *Culture*, where I include writers, painters, composers, etc. The third group called *Athletes and entertainers* consists mostly of athletes but also includes actors and singers. The exact list is in Appendix B in Table B.1.

As of July 2014, 259,638 registered editors of English-language Wikipedia had reported their gender when registering their Wikipedia username. Of those, 14.4 percent reported their gender as female. The number of editors who reported their gender is not small; however, the majority of registered editors have not reported their gender. This is not surprising because Wikipedia requires almost no information to register, no real name or even an email address.

The editors who report their gender tend to be more active, with a larger number of edits. Note that for the analysis here it is not important that male and female editors report their gender at equal rates. For example, it might be that women are less likely to report their



gender for fear of harassment. For this analysis that is not a problem. What is important is that editors don't systematically report the wrong gender. Since the default option is not specifying one's gender, I would not expect that they are massively reporting the wrong gender.

In this appendix, I analyze the behavior of all the editors who reported their gender and who have edited English-language Wikipedia articles about human beings in the Wikidata sample. There are 74,817 such editors. They have edited 866,431 articles about the human beings in the Wikidata sample.

**Summary statistics** Table A.1 presents summary statistics of the Wikipedia editors. 12.6 percent of editors are women. Women contribute somewhat longer texts than men, both per article and in total per year.

Table A.1: Summary statistics of Wikipedia editors, split by gender

Percentiles	Male editors			Female editors		
	25th	50th	75th	25th	50th	75th
Number of biographies edited, per year	1	2	4	1	1	3
Length of text added, per year	22	169	1169	26	188	1247
Length of text added, per biography	13	80	294	16	98	474
Observations	65398			9419		

Note: A unit of observation is an editor. Columns 1-3, present the 25th, 50th, and 75th percentile for male editors, and columns 4 - 6, for female editors. *Length of text added* measures in characters the length of text added to biographies in Wikipedia. Sample of editors includes all the editors who reported their gender and who have edited the English language Wikipedia articles about human beings in the Wikidata final sample.

Changes over time in the number of active editors and the share of active female editors are depicted in Figure A.1. By active, I mean that the editor has edited at least one article during a given year. In 2013, 10.4 percent of the active editors were female. The percentage of female editors among active editors has increased from 3.7 percent in 2002 to a peak of 11.5 percent in 2011. Changes in the percentage of female editors mirrored changes in the total number of active editors. The number of active editors increased monotonically from 81 in 2002 to 27,600 in 2011 and was 22,600 in 2013.

## A.2 The effect of the gender gap on Wikipedia content: evidence from Wikipedia editing histories

In this part, I simply document that Wikipedia female editors direct relatively more of their editorial input to biographies of women. Table A.2 presents the average (over editors)




Figure A.1: Number of active Wikipedia editors (left y-axes) and percentage of active female editors (right y-axes), in 2002 - 2013

Note: The dataset is restricted to the registered editors of Wikipedia who have reported their gender and edited at least one article in a given calendar year.

percentage of contributions to biographies of women from contributions to all biographies. An observation here is an editor. In column 2, the average percentage is calculated over male editors and, in column 3, over female editors. Contributions can be measured in multiple ways. In rows 1-3, contributions are measured in the length of text added.

The first row of Table A.2 shows that the average percentage of text added to biographies of women is 31% in the case of female editors and 17% in the case of male editors. To make sure that the gender difference does not come only from celebrities like actors, athletes, and pop stars, rows 2 and 3 restrict attention to subgroups of biographies. In row 2, the sample includes lawyers, politicians, scientists, etc. (group Professionals); in row 3, it includes writers, painters, composers, etc. (group Culture); the exact list of the groups is presented in Appendix B, Table B.1. A large gender difference remains in both groups. Rows 4-6 measure contributions by the number of biographies edited. The measure gives similar results.

Figure A.2 presents the percentage of text added to biographies of women from text added to all biographies, separately by male and female editors. Excluding the beginning of Wikipedia when the number of editors was very small, the relative share of contributions to biographies of women has increased, but the difference between male and female editors remained large.

Figure A.3 presents the same statistics for the same subgroups of biographies as above. The trends are similar to aggregate statistics.

Table A.2: The average (over editors) percentage of contributions to the biographies of women vs men, by the gender of Wikipedia editor

	% biographies of women		t-test	Wilcoxon test	Obs.
	Men	Women	p-value	p-value	
Total length of text added	16.6	31.4	0.000	0.000	66112
– Subset of biographies: Professionals	8.0	16.7	0.000	0.000	26459
– Subset of biographies: Culture	15.7	31.7	0.000	0.000	22665
Number of biographies edited	17.1	31.1	0.000	0.000	74817
– Subset of biographies: Professionals	8.3	15.9	0.000	0.000	30635
– Subset of biographies: Culture	15.7	31.5	0.000	0.000	26360

Note: A unit of observation is a Wikipedia editor. In rows 1 - 3, the average percentage of contributions to the biographies of women is measured in terms of the total length of text added. In Columns 1 and 2, present the means separately for male and female Wikipedia editors. Column 3, presents the p-value of the t-test for whether the difference between male and female editors is significantly different from zero. Column 4, presents the p-value of the corresponding Wilcoxon rank-sum test. Column 5 presents the number of observations used in each test. In rows 1 - 3, the average percentage of contributions to the biographies of women is measured in terms of the total length of text added. In rows 4 - 6, it is measured in terms of the number of biographies edited. Sample includes 65,398 male and 9,419 female editors. In rows 1–3, the sample is restricted to editors who have added (instead of only deleted) text to biographies. In rows 2 and 5, the sample is restricted to the biographies of lawyers, politicians, scientists, etc (*Professionals*). In rows 3 and 6, the sample is restricted to the biographies of writers, painters, composers, etc (*Culture*). The exact definition of these groups of professions is in Table B.1.

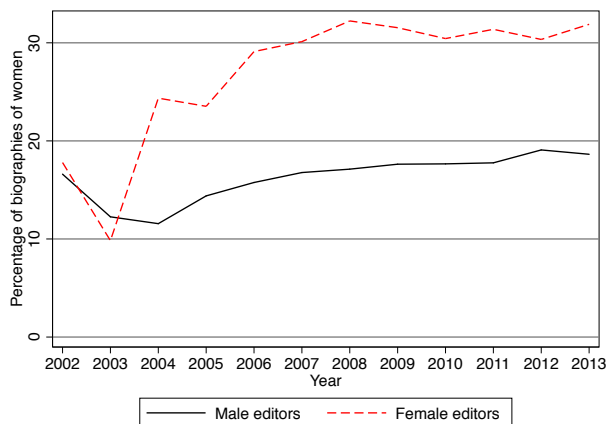


Figure A.2: Average (over Wikipedia editors) percentage of contributions to the biographies of women vs men, in 2002 - 2013

Note: The percentage is calculated as the average (over Wikipedia editors) percentage of the total length of text added to the biographies of women vs men.

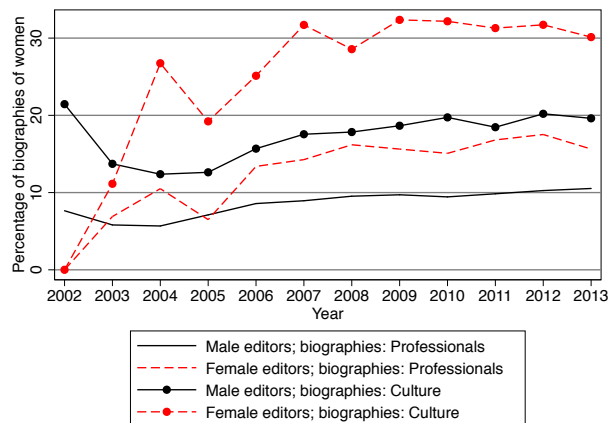


Figure A.3: Average (over Wikipedia editors) percentage of contributions to the biographies of women vs men, by gender of editor and profession of the person in the biography, in 2002 – 2013

Note: The percentage is calculated as the average (over editors) percentage of the total length of text added to the biographies of women in the total length added to all biographies. The sample of *Male editors; biographies: Professionals* is restricted to male editors and a subgroup of biographies that includes lawyers, politicians, scientists, etc; the sample of *biographies: Culture* includes writers, painters, composers, etc; the exact definition of the subgroups of biographies is in Table B.1.

### A.3 Is the allocation of editorial input proportional to the readership?

This section analyzes whether the allocation of editorial input to the Wikipedia articles about men and women is proportional to the readership. Each day, there is a large share of Wikipedia articles that no one reads. I will start by restricting attention to those. Second, I will look at all the biographies.

I use data on Wikipedia viewership statistics. The dataset includes for each Wikipedia page the number of page views in September 2014 for more than 900,000 biographies in Wikipedia..<sup>33</sup> Page views are from a time period outside the editing histories in the sample. On average, the number of page views is rather stable across months, and there is no reason to believe that September 2014 is somehow different from other time periods.

First, I simply document that on a typical day, the percentage of biographies of men that no one reads is larger than that of women. Table A.3 presents the daily statistics regarding the percentage of biographies in Wikipedia that no one reads. On a typical (median) day in September 2014, no one read 26 percent of the biographies of men versus only 16 percent of the biographies of women.<sup>34</sup> Note that over a longer period of time almost all the articles

<sup>33</sup> The data on page views is downloaded from: <http://dumps.wikimedia.org/other/pagecounts-raw/>

<sup>34</sup>The median is taken over the days, and it is calculated separately for the biographies of men and the

have readers. The above statistics could be interpreted that at the lower tail of the readership distribution, biographies of men have fewer readers than biographies of women. We see this also when we look at the whole distribution of readership.

Table A.3: The daily percentage of biographies in Wikipedia that noone read

	Median	Min	Max	Obs
Biographies of men	25.8	18.6	28.4	30
Biographies of women	16.2	11.9	18.6	30

Note: A unit of observation is a day. The sample consists of daily data of page views from September 2014 for 784,178 biographies of men and 140,193 biographies of women in Wikipedia.

Figure A.4 presents the cumulative distribution of the number of page views of biographies of men and women, where the page views are divided by the number of editors in the sample who edited the page. As confirmed by the Kolmogorov-Smirnov test in Table A.4, biographies of women receive more page views per editor.

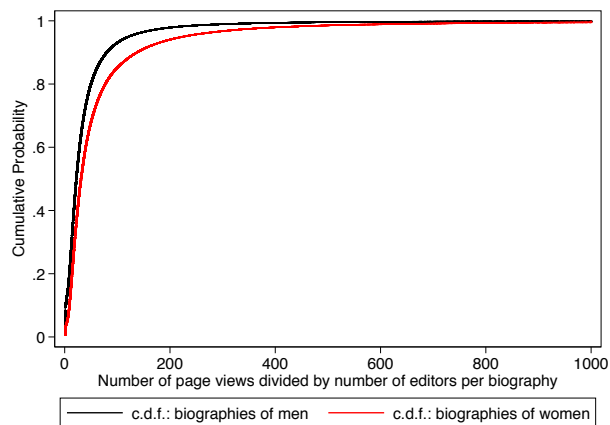


Figure A.4: Cumulative distribution functions of the number of page views divided by the number of editors per biography

Note: The sample includes page views data from September 2014 for 734,767 biographies of men and 131,664 biographies of women in Wikipedia. The number of page views divided by the number of editors is constructed based on the pages and editors in the sample of historical editing data from Wikipedia. The sample excludes biographies that the editors in the sample did not edit.

---

biographies of women.

Table A.4: Kolmogorov-Smirnov test of whether the samples of the number of page views divided by the number of editors of the pages of men and women are drawn from the same distribution

	Test statistic	p-value	Observations
One-sided hypothesis test	0.1414	0.000	866431
Two-sided hypothesis test	0.1414	0.000	866431

Note: The number of page views divided by the number of editors is constructed based on the pages and editors in the sample of historical editing data from Wikipedia. Page views data from September 1, 2014. The null hypothesis is that the samples are drawn from the same distribution. In row 1, the alternative hypothesis is that the c.d.f. of the pages of women is larger than the c.d.f. of the pages of men. In row 2, the alternative hypothesis is that the samples are drawn from two different distributions. The sample excludes biographies that the editors in the sample did not edit.

Appendix E.1 repeats the analysis for survey data, and the results are similar. Overall, this seems to imply that demand side cannot be the only reason for gender inequality in Wikipedia coverage. Instead, gender inequality in coverage seems to originate, at least partially, from the supply side of Wikipedia production.

## B Online Appendix: Additional tables and figures

Imagine a situation where you were asked to improve Wikipedia. Specifically, suppose you were asked to do the following:

1. Choose a Wikipedia page of any **human being** (it could be a scientist, author, journalist, politician, etc.; either living or dead; exclude fictional characters).
2. Find some relevant information from the web that is currently missing from the Wikipedia page, and add that information to the page (at least one sentence). The information added should be based on a previously published source (article, book, newspaper, etc.). When sources disagree, present what the various sources say.

Please take some time to consider the hypothetical situation, look at a few Wikipedia articles, and search for the relevant information on the web. Then, please, answer below, which **human being** would you choose, and what would you add to the page.

**\* Which human being would you choose? Please write here the name of the human being.**

**\* Please provide a link to the Wikipedia page of the human being, the page that you would choose to edit.**  
Please make sure that it is a valid link to the page of the human being on English language Wikipedia; it must start with:  
<http://en.wikipedia.org/wiki/>

**\* Please write here the text that you would add to the page.**

Figure B.1: Survey instructions of Wikipedia hypothetical editing

Table B.1: Professions of human beings in Wikidata

Profession	Number of people	Category
football player	105871	Athletes & entertainers
actor	42096	Athletes & entertainers
politician	40084	Professionals
sportsperson	23040	Athletes & entertainers
baseball player	20423	Athletes & entertainers
writer	15515	Arts & literature
cricketer	15448	Athletes & entertainers
lawyer	14812	Professionals
painter	12746	Arts & literature
ice hockey player	11751	Athletes & entertainers
composer	11692	Arts & literature
priest	11242	Professionals
film director	9794	Athletes & entertainers
basketball player	8561	Athletes & entertainers
singer	8122	Athletes & entertainers
journalist	7619	Professionals
poet	7273	Arts & literature
screenwriter	6807	Athletes & entertainers
musician	6325	Athletes & entertainers
diplomat	6214	Professionals
Australian-rules footballer	6122	Athletes & entertainers
judge	5767	Professionals
officer	5509	Professionals
mathematician	5352	Professionals
linguist	5229	Professionals
photographer	4750	Arts & literature
tennis player	4498	Athletes & entertainers
conductor	4428	Arts & literature
boxer	4305	Athletes & entertainers
rugby league player	4222	Athletes & entertainers
physicist	4197	Professionals
film producer	4192	Athletes & entertainers
economist	3922	Professionals
author	3485	Arts & literature
architect	3425	Professionals
anthropologist	3082	Professionals
bicycle racer	3065	Athletes & entertainers
basketball coach	3048	Athletes & entertainers
golfer	2919	Athletes & entertainers

Table continues on next page



Profession	Number of people	Category
explorer	2856	Professionals
singer-songwriter	2841	Athletes & entertainers
chemist	2651	Professionals
botanist	2585	Professionals
film actor	2432	Athletes & entertainers
rugby union player	2374	Athletes & entertainers
astronomer	2271	Professionals
television actor	2164	Athletes & entertainers
motorcycle racer	1961	Athletes & entertainers
television presenter	1924	Athletes & entertainers
sculptor	1915	Arts & literature
historian	1904	Professionals
theologian	1672	Professionals
model	1606	Athletes & entertainers
philosopher	1594	Professionals
biologist	1580	Professionals
computer scientist	1546	Professionals
physician	1515	Professionals
pianist	1457	Arts & literature
field hockey player	1452	Athletes & entertainers
voice actor	1409	Athletes & entertainers
entrepreneur	1282	Professionals
alpine skier	1281	Athletes & entertainers
swimmer	1255	Athletes & entertainers
librarian	1211	Professionals
soldier	1159	Professionals
civil engineer	1075	Professionals
badminton player	1064	Athletes & entertainers

Table continued from previous page

Table B.2: Estimated role of survey respondent’s characteristics for Wikipedia editing, by gender

	Has edited Wikipedia		Likely to edit	
	Men (1)	Women (2)	Men (3)	Women (4)
At least 35 years old	-0.078 (0.061)	0.038 (0.048)	-0.049 (0.046)	0.005 (0.039)
College degree	0.127** (0.055)	-0.001 (0.049)	0.050 (0.043)	0.003 (0.040)
Unemployed	-0.060 (0.112)	0.019 (0.072)	-0.046 (0.119)	0.043 (0.060)
Works at least 35 hours per week	-0.095 (0.062)	-0.014 (0.049)	-0.044 (0.047)	0.066 (0.042)
Use Wikipedia daily	0.170*** (0.055)	0.148*** (0.052)	0.127*** (0.042)	0.080* (0.044)
Use Wikipedia at least weekly	0.079 (0.091)	0.059 (0.066)	0.076 (0.086)	-0.017 (0.049)
Believe: more competent than other editors	0.146 (0.091)	0.109 (0.095)	0.052 (0.058)	-0.044 (0.089)
Believe: at least as competent as other editors	0.076 (0.071)	0.150** (0.062)	0.266** (0.111)	0.122** (0.061)
Donated to charity	0.042 (0.061)	0.041 (0.054)	0.078 (0.048)	0.010 (0.042)
Volunteering at least one hour per week	-0.022 (0.063)	-0.030 (0.050)	0.008 (0.047)	0.047 (0.039)
Wrote user review in previous six months	0.089 (0.055)	0.098** (0.049)	0.045 (0.043)	0.012 (0.038)
Leisure at least 3 hours per weekday	0.003 (0.062)	-0.031 (0.054)	-0.025 (0.044)	0.056 (0.048)
Log-likelihood	-144.593	-90.465	-87.920	-63.547
Observations	259	233	259	233

Note: Each column presents estimated marginal effects from a separate logit regression. In columns 1 and 2, the dependent variable is an indicator of whether the individual has edited Wikipedia in the past. In columns 3 and 4, the dependent variable is whether the individual is likely to edit Wikipedia in the next 30 days, specifically, the variable takes value one if the respondent answered that he is either “very likely” or “quite likely” to the question: “How likely or unlikely is it that you will actually edit Wikipedia over the next 30 days?” The sample in columns 1 and 3 is restricted to men in the control group; in columns 2 and 4, it’s women in the control group. *Believe: more competent than other editors* is an indicator variable that takes value 1 if answered “More knowledgeable/competent” to the question “Think of the Wikipedia article of the human being that you chose. On the whole, do you think that you are more or less knowledgeable and competent to edit the article than other people who will edit it in the future?” *Believe: at least as competent as other editors* is an indicator variable that takes value 1 if answered either “More knowledgeable/competent”, “Somewhat more”, or “Same” to the same question above. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

Table B.3: Differences in the effect of treatment on the likelihood of editing Wikipedia: by Wikipedia use, believed competence, and whether writes user reviews

	Wikipedia use		Competence		Writing reviews	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	-0.048**	-0.049**	-0.064**	-0.066**	-0.058**	-0.058**
	(0.022)	(0.022)	(0.026)	(0.026)	(0.028)	(0.029)
Covariates	No	Yes	No	Yes	No	Yes
Control group mean	0.136	0.136	0.161	0.161	0.146	0.146
Scaled treatment effect	-0.353	-0.357	-0.395	-0.406	-0.396	-0.398
Observations	788	788	691	691	511	511

Note: Each column presents estimates from a separate regression. Dependent variable is an indicator for whether the individual believes he is likely to edit Wikipedia in the next 30 days, defined in the same way as in Table 2. Covariates include demographic characteristics from Table 2. The sample includes both the control and treatment group of the survey respondents. In columns 1-2, the sample is restricted to those who use Wikipedia at least weekly. In columns 3-4 the sample is restricted to those who believe they are at least as competent and knowledgeable as other people who will edit Wikipedia in the future. In columns 5-6 the sample is restricted to those who wrote a user review during the past six months. *Scaled treatment effect* measures the percentage change of the share of respondents who believe it is likely they will edit Wikipedia. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

## C Online Appendix: Survey attrition

The survey had an overall attrition rate of 19%.<sup>35</sup> This section tests whether attrition is random or depends on demographic characteristics and treatment status.

98% of the individuals who dropped out, did that on the page of the survey that displayed the task of choosing the Wikipedia page which to edit and finding information what to add to the page. This is probably where the participants realized that the task is more difficult than they expected. Two percent of those who dropped out did that before seeing the treatment page. None dropped out on the treatment page. Altogether, this is likely to indicate that treatment itself was not a reason for dropping out.

Table C.1 shows that although attrition is not random, it doesn't depend on treatment status nor gender. However, as could be expected, people who use or edit Wikipedia, are more likely to finish the survey. Table C.1 presents estimates from linear regressions<sup>36</sup> of the form:  $FinishedSurvey_i = Constant + \beta Variable_i + \varepsilon_i$ , where the dependent variable  $FinishedSurvey_i$  is an indicator of whether person  $i$  finished the survey. Each row of the table presents the estimate of coefficient  $\beta$  and corresponding p-value from a regression where the  $Variable_i$  of the regression is listed in the first column. In each regression, the sample includes only those respondents who have remained in the survey to answer the question; in case of treatment status, the sample includes those who have remained in the survey until they were assigned the status. The first row of the table shows that gender has no statistically significant impact on finishing the survey. The last row shows that treatment status has no statistically significant impact either. The table also shows that attrition depends on demographic characteristics and those who use Wikipedia more are more likely to finish the survey. Specifically, people who have at least a college degree are more likely to finish the survey; and those who are older or unemployed are less likely to finish the survey. People who use Wikipedia at least weekly or who have edited it, are more likely to finish the survey.

To test it further that the attrition does not depend on treatment status, Table C.2 presents estimated marginal effects from logit regressions, where the dependent variable is an indicator for finishing the survey. It is regressed on the treatment status and demographic characteristics (in columns 1-2) and treatment status interacted with the demographic characteristics (in column 2). Treatment status has no significant effect on finishing the survey. According to the specification in column 2, there are four characteristics that matter for attrition: those who have edited Wikipedia before are 13 percentage points more likely to finish the survey,

---

<sup>35</sup> When counting the number of people who dropped out, I exclude 2 responses that don't have unique Amazon Mechanical Turk IDs. When calculating the attrition rate, I include only the 972 valid completed responses and the uncompleted responses.

<sup>36</sup>Estimates from analogous logit regressions are available from the author. The results are very similar.

Table C.1: Ability of covariates to predict whether respondents finish the survey

	Coefficient	p-value	Observations
Female	-0.026	0.252	1203
At least 35 years old	-0.058	0.014	1203
Has college degree	0.045	0.050	1203
Works at least 35 hours per week	0.036	0.112	1203
Unemployed	-0.097	0.007	1203
Use Wikipedia daily	0.026	0.307	1202
Use Wikipedia at least weekly	0.056	0.046	1202
Has edited Wikipedia	0.117	0.000	1202
Treatment group	0.028	0.210	1202

Note: Each row presents estimates from a separate linear regression of the form  $FinishedSurvey_i = Constant + \beta Variable_i + \varepsilon_i$ , where the  $Variable_i$  is listed in the first column. A unit of observation is a survey respondent. In each regression, the sample includes only those respondents who have remained in the survey to answer the question; in case of treatment status, the sample includes those who have remained in the survey until they were assigned the status. Number of observations is 1203 for demographic characteristics and 1202 for Wikipedia related questions and treatment status.

those who have a college degree are 5 percentage points more likely to finish it, and those who are older (at least 35 years old) or unemployed are less likely to finish it.

Table C.3 shows that conditional on finishing the survey, assignment into a treatment group was random in terms of most characteristics. The table presents estimates from linear regressions<sup>37</sup> of the form:  $TreatmentGroup_i = Constant + \beta Variable_i + \varepsilon_i$ , where the dependent variable  $TreatmentGroup_i$  equals one if person  $i$  was assigned into the treatment group and zero if he was assigned into the control group. Each row of the table presents the estimate of coefficient  $\beta$  and corresponding p-value from a separate regression where the  $Variable_i$  of the regression is listed in the first column. In each regression, the sample includes only those respondents who finished the survey. The table shows that conditional on finishing the survey, assignment into treatment group was random in terms of the demographic characteristics, attitudes towards Wikipedia, the frequency of Wikipedia use, charity, writing user reviews, and free time. From the 14 outcomes, one is significant at the 10 percent level. Those who have edited Wikipedia before, conditional on finishing the survey are less likely to be in the treatment group, with p-value equal to 0.08. Note that the question was asked before treatment, hence, the treatment couldn't affect the answers.

The treatment provided information about Wikipedia and was designed to change Wikipedia editing behavior. Two concerns were that the treatment could lead to differential attrition or change the answers to the survey questions. From Table C.1 we saw that

<sup>37</sup>Estimates from analogous logit regressions are available from the author. The results are very similar.

Table C.2: Ability of covariates to predict whether respondents finish the survey

	(1)	(2)
Female	-0.002 (0.023)	-0.027 (0.032)
At least 35 years old	-0.048** (0.023)	-0.066** (0.031)
Has college degree	0.030 (0.023)	0.051* (0.031)
Works at least 35 hours per week	0.005 (0.025)	-0.013 (0.035)
Unemployed	-0.078** (0.035)	-0.085* (0.047)
Use Wikipedia daily	-0.012 (0.028)	-0.031 (0.038)
Use Wikipedia at least weekly	0.031 (0.028)	0.010 (0.038)
Has edited Wikipedia	0.130*** (0.036)	0.132*** (0.047)
Treatment group	0.030 (0.022)	-0.065 (0.070)
Treatment group * Female		0.051 (0.047)
Treatment group * At least 35 years old		0.042 (0.046)
Treatment group * Has college degree		-0.045 (0.046)
Treatment group * Works at least 35 hours per week		0.040 (0.051)
Treatment group * Unemployed		0.010 (0.070)
Treatment group * Use Wikipedia daily		0.042 (0.056)
Treatment group * Use Wikipedia at least weekly		0.051 (0.057)
Treatment group * Has edited Wikipedia		-0.003 (0.074)
Log-likelihood	-568.966	-566.551
Observations	1202	1202

Note: Each column presents estimated marginal effects from a separate logit regression. Dependent variable is an indicator of whether the respondent finished the survey. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

there is no evidence that the treatment lead to differential attrition. Table C.3 showed that in terms of most characteristics the randomization into treatment was successful. In any case, when analyzing the survey responses, I restrict attention to the control group, to make sure that the treatment does not affect the answers.

Table C.3: Ability of covariates to predict treatment status, conditional on finishing the survey

	Coefficient	p-value	Observations
Female	0.043	0.180	972
At least 35 years old	0.012	0.725	972
Has college degree	0.030	0.352	972
Works at least 35 hours per week	0.034	0.302	972
Unemployed	-0.016	0.762	972
Use Wikipedia daily	0.011	0.757	972
Use Wikipedia at least weekly	0.059	0.147	972
Has edited Wikipedia	-0.067	0.083	972
Believe: more competent than other editors	-0.002	0.977	972
Believe: at least as competent as other editors	-0.016	0.648	972
Donated to charity	-0.035	0.274	972
Volunteering at least one hour per week	-0.018	0.592	972
Wrote user review in previous six months	-0.010	0.763	972
Leisure at least 3 hours in a weekday	0.016	0.664	972

Note: Each row presents estimates from a separate regression of the form  $TreatmentGroup_i = Constant + \beta Variable_i + \varepsilon_i$ , where  $TreatmentGroup_i$  equals one if person  $i$  was assigned into the treatment group and zero if he was assigned into the control group, the  $Variable_i$  is listed in the first column. A unit of observation is a survey respondent. In each regression, the sample includes only those 972 respondents who have finished the survey.

## D Online Appendix: Analysis with the restricted sample

Table D.1: Repeating Table 5 with the restricted sample. Likelihood of choosing to edit a Wikipedia article about a woman vs a man, by the gender of survey respondent

	Chose an article about a woman		t-test	Obs.
	Male respondents	Female respondents	p-value	
All respondents in the control group				
Total	0.115	0.317	0.000	417
Subsets of respondents by demographic characteristics				
Has edited Wikipedia	0.154	0.382	0.010	99
Hasn't edited Wikipedia	0.098	0.303	0.000	318
College	0.121	0.303	0.001	225
No college	0.108	0.333	0.000	192
Less than 35 years of age	0.069	0.284	0.001	146
At least 35 years of age	0.137	0.336	0.000	271
Subsets of respondents by type of Wikipedia articles				
Biographies: Professionals	0.095	0.217	0.049	134
Biographies: Culture	0.085	0.267	0.007	134

Note: A unit of observation is a survey respondent. Each survey respondent chose a Wikipedia article either about a man or a woman. Columns 1 and 2 present the mean of the indicator variable that the chosen article was about a woman, separately for male and female survey respondents. Column 3 presents the p-value of the t-test for whether the difference between male and female respondents is significantly different from zero. In the first row, the sample consists of the survey respondents in the control group whose answers were categorized as providing information as described in subsection 3.2. In the following rows, the sample is further restricted by the survey respondents' demographic characteristics that are listed in the leftmost column. In the last two rows, the sample is restricted to the subgroups of biographies: group *Professionals* includes lawyers, politicians, scientists, etc; and group *Culture* includes writers, painters, composers, etc; the exact definition of groups is in Table B.1.



Table D.2: Repeating Table 6 with the restricted sample. Estimated role of survey respondents' characteristics on the likelihood of choosing to edit a Wikipedia article about a woman vs a man

	Chose an article about a woman	
	Coefficients (1)	Marginal effects (2)
Female	1.361*** (0.280)	0.210*** (0.040)
At least 35 years old	-0.393 (0.274)	-0.061 (0.042)
College degree	-0.056 (0.261)	-0.009 (0.040)
Unemployed	-0.279 (0.444)	-0.043 (0.068)
Works at least 35 hours per week	-0.106 (0.275)	-0.016 (0.042)
Use Wikipedia daily	0.084 (0.322)	0.013 (0.050)
Use Wikipedia at least weekly	-0.039 (0.322)	-0.006 (0.050)
Has edited Wikipedia	0.426 (0.305)	0.066 (0.047)
Log-likelihood	-199.394	
Observations	417	

Note: Estimation results are from a logit regression. Column 1 presents coefficient estimates and column 2 presents estimated marginal effects. Dependent variable is an indicator for whether the respondent chose to edit an article about a woman. The explanatory variables are the characteristics of the survey respondent. The sample is restricted to the survey respondents in the control group whose answers were categorized as providing information as described in subsection 3.2. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

Table D.3: Repeating Table 7 with the restricted sample. The effect of treatment on the likelihood of choosing to edit a Wikipedia article about a woman

	Total		Male respondents		Female respondents	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.195***	0.187***	0.130***	0.132***	0.247***	0.250***
	(0.031)	(0.030)	(0.037)	(0.037)	(0.048)	(0.048)
Covariates	No	Yes	No	Yes	No	Yes
Control group mean	0.211	0.211	0.115	0.115	0.317	0.317
Scaled treatment effect	0.926	0.888	1.136	1.149	0.779	0.788
Observations	823	823	418	418	405	405

Note: Each column presents estimates from a separate linear probability model. Dependent variable is an indicator for whether the respondent chose to edit an article about a woman. Covariates include demographic characteristics from Table 2. The sample includes both the control and treatment group of the survey respondents but is restricted to those whose answers were categorized as providing information as described in subsection 3.2. The sample is restricted to men in columns 3-4 and women in columns 5-6. *Scaled treatment effect* measures the percentage change of the share of women’s pages. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

Table D.4: Repeating Table 8 with the restricted sample. The effect of treatment on the intention to edit Wikipedia in the future

	Total		Male respondents		Female respondents	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	-0.039*	-0.041**	-0.078***	-0.081***	0.002	-0.002
	(0.021)	(0.021)	(0.029)	(0.029)	(0.029)	(0.030)
Covariates	No	Yes	No	Yes	No	Yes
Control group mean	0.115	0.115	0.133	0.133	0.095	0.095
Scaled treatment effect	-0.337	-0.358	-0.587	-0.611	0.017	-0.025
Observations	823	823	418	418	405	405

Note: Each column presents estimates from a separate linear probability model. Dependent variable is an indicator for whether the individual believes he is likely to edit Wikipedia in the next 30 days, defined in the same way as in Table 2. Covariates include demographic characteristics from Table 2. The sample includes both the control and treatment group of the survey respondents but is restricted to those whose answers were categorized as providing information as described in subsection 3.2. The sample is restricted to men in columns 3-4 and women in columns 5-6. *Scaled treatment effect* measures the percentage change of the share of respondents who believe it is likely they will edit Wikipedia. Standard errors are in parentheses. \*\*\* Significant at the 1 percent level. \*\* Significant at the 5 percent level. \* Significant at the 10 percent level.

## E Online Appendix: Additional empirical analysis

### E.1 Is the allocation of editorial input proportional to the readership: evidence from the survey

The evidence from the Wikipedia editing histories showed that pages of men receive relatively fewer readers than the pages of women. A similar exercise with the survey data confirms the results. The analysis uses survey respondents stated choices of the pages they would choose to edit.

Figure E.1 presents the cumulative distribution of the number of page views of the Wikipedia pages of men and women, where the page views are divided by the number of survey respondents editing the page. The measure is constructed only based on the editors and pages from the control group in the survey.

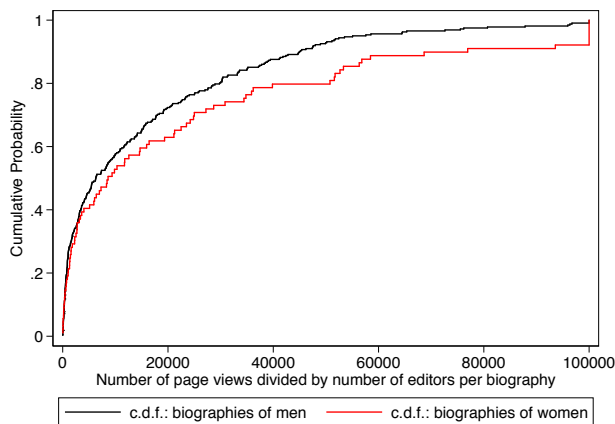


Figure E.1: Cumulative distribution functions of the number of page views divided by the number of editors from the survey editing the article, by gender (of the person in the article)

Note: Page views data is from September 2014 for the first week after the survey was conducted. The number of page views divided by the number of editors is constructed based on the editors and articles from the control group in the survey.

As partly confirmed by the Kolmogorov-Smirnov test, presented in Table E.1, the c.d.f. of the pages of women seems to be larger than that of the pages of men. In Table E.1, the null hypothesis is that the samples are drawn from the same distribution. The first row of the table presents the test statistic and p-value from the one-sided hypothesis test, where alternative hypothesis is that the c.d.f. of the pages of women is larger than the c.d.f. of the pages of men. The one-sided test rejects the null hypothesis at the 10 percent significance level. The second row of the table presents the test statistic and p-value from the two-sided hypothesis test, where the alternative hypothesis is that the samples are drawn from two

different distributions. The two-sided test does not reject the null hypothesis.

Table E.1: Kolmogorov-Smirnov test of whether the samples of the number of page views divided by the number of editors of the pages of men and women are drawn from the same distribution

	Test statistic	p-value	Observations
One-sided hypothesis test	0.1339	0.082	411
Two-sided hypothesis test	0.1339	0.147	411

Note: A unit of observation is a Wikipedia page. The sample is restricted to pages that were chosen by the survey respondents in the control group. The number of unique pages is smaller than the number of survey respondents because some pages were chosen more than once. The null hypothesis is that the samples are drawn from the same distribution. In row 1, the alternative hypothesis is that the c.d.f. of the pages of women is larger than the c.d.f. of the pages of men. In row 2, the alternative hypothesis is that the samples are drawn from two different distributions.

Past literature has shown that images drive viewership in Wikipedia. Table E.2 presents results from a regression where the dependent variable is the logarithm of page views and explanatory variables are the number of images on the page and an indicator whether it is a biography of a woman. The analysis is subject to the caveat that the data on images is from a later period (January 2019) than the data on page views (September 2014). The results show that even controlling for the number of images on the page, biographies of women still have more views.

Table E.2: Dependent variable: logarithm of page views

Number of images	0.089*** (0.007)
Biography of a woman	0.541** (0.212)
Constant	8.070*** (0.119)
Observations	492
Adj. R-squared	0.268

Note: A unit of observation is a survey respondent. The sample is restricted to the survey respondents in the control group. Dependent variable is the logarithm of views of the survey respondent's chosen Wikipedia page. Page views data is from September 2014 for the first week after the survey was conducted. Data on images is from January 2019.

## E.2 In terms of total readership, does the treatment improve the allocation of editorial input?

In this section, I compare the allocation of editorial input in the treatment and control group. The goal is to learn whether in the treatment group compared to the control group more editorial input is allocated to the pages with a larger or smaller number of readers. The analysis uses survey respondents stated choices of the pages they would choose to edit.

Figure E.2 presents the cumulative distribution of the number of page views of the Wikipedia pages in the control and treatment groups, where the page views are divided by the number of survey respondents editing the page in control and treatment groups respectively. As confirmed by the Kolmogorov-Smirnov test in Table E.3, the distributions look similar. In Table E.3, the null hypothesis is that the samples are drawn from the same distribution. The first row of the table presents the test statistic and p-value from the one-sided hypothesis test, where alternative hypothesis is that the c.d.f. of the pages in the control group is smaller than the c.d.f. of the pages in the treatment group. The second row presents the results from the one-sided hypothesis test, where the alternative hypothesis is that the c.d.f. of the pages in the treatment group is smaller than the c.d.f. of the pages in the control group. The third row presents results from the two-sided hypothesis test, where the alternative hypothesis is that the samples are drawn from two different distributions. All three tests cannot reject the null hypothesis that the samples are drawn from the same distribution at the 10 percent significance level.

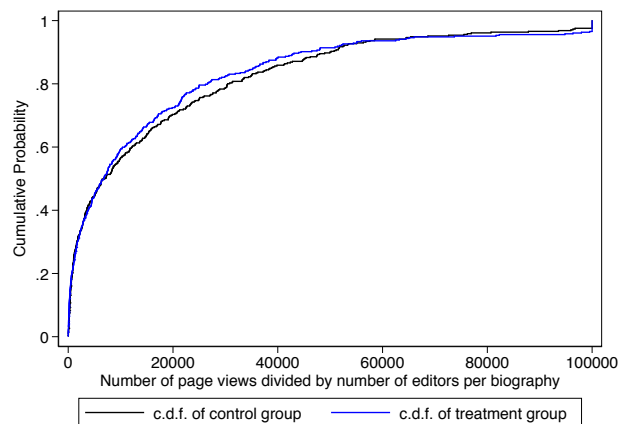


Figure E.2: Cumulative distribution functions of the number of page views divided by the number of editors from the survey editing the article, by control and treatment group

Note: Page views data is from September 2014 for the first week after the survey was conducted. The number of page views divided by the number of editors is constructed based on the editors and articles from the survey.

Hence, we can conclude that although the treatment moved the allocation of editorial input towards the pages of women, this didn't make the allocation worse nor better in terms of the readership.

Table E.3: Kolmogorov-Smirnov test of whether the samples of the number of page views divided by the number of editors of the pages in control and treatment groups are drawn from the same distribution

	Test statistic	p-value	Observations
One-sided hypothesis test: control group smaller	0.0269	0.744	818
One-sided hypothesis test: treatment group smaller	-0.0468	0.408	818
Two-sided hypothesis test	0.0468	0.735	818

Note: A unit of observation is a Wikipedia page. The sample includes pages chosen by the survey respondents either in the treatment or control group. The number of unique pages is smaller than the number of survey respondents because some pages were chosen more than once. The null hypothesis is that the samples are drawn from the same distribution. In row 1, the alternative hypothesis is that the c.d.f. of the pageviews of the pages chosen in the control group are smaller than the c.d.f. of the pages in the treatment group. In row 3, the alternative hypothesis is that the samples are drawn from two different distributions.