GENDER BIAS IN RUMORS AMONG PROFESSIONALS: AN IDENTITY-BASED INTERPRETATION

Alice H. Wu*

Abstract—This paper measures gender bias in discussions about women versus men in an online professional forum. I study the content of posts that refer to each gender and the transitions in the topics between consecutive posts once attention turns to one gender or the other. Discussions about women tend to emphasize their personal characteristics instead of professional accomplishments. Posts about women are also more likely to lead to deviations from professional topics than are posts about men. I interpret these findings through a model that highlights posters' incentives to boost their own identities relative to the underrepresented out-group in a profession.

I. Introduction

OCCUPATIONAL segregation by gender has been declining but at a slower pace in the past decades (Blau, Brummund, & Liu, 2013). Gender gaps also persist in math-intensive fields like economics, engineering, and computer science (Ceci et al., 2014), and institutional efforts to promote integration in these fields often face a backlash, as evidenced by the highly publicized antidiversity memo from a software engineer at Google. Some analysts argue that gender role attitudes have changed little since the 1990s, and discrimination at the workplace today can take a subtler form than blatant expressions of sexism (Cotter, Hermsen, & Vanneman, 2011; Basford, Offermann, & Behrend, 2013). In particular, the increasing share of women may still be perceived by men as diluting (or "polluting") the rigor and prestige of a profession (Goldin, 2015).

Understanding attitudes toward gender among colleagues is important because such attitudes may contribute to a stereotypical professional climate that discourages women from entering and staying in certain fields and leads to a persistent underrepresentation of women. Yet it remains challenging to study this issue in real-world settings where people who are concerned about social correctness will not readily reveal their beliefs about gender.

This paper aims to measure gender bias in an anonymous online setting where members of the economics profession are presumably freed from social pressure and thus are more likely to reveal their true gender attitudes. Economics is one of the largest academic disciplines where men still substantially outnumber women at both student and faculty levels (Lundberg, 2018). The persistently low share of women has attracted substantial interest and concern (see Bayer & Rouse, 2016, for a summary), and recent research on publications, a key performance metric for economists, suggests that women face a higher bar than men in the peer review process and are given less credit when collaborating with men (Card et al., 2019; Hengel, 2019; Sarsons et al., forthcoming). A new professional climate survey conducted by the American Economic Association also finds that women are much less likely to feel included socially or intellectually within economics and more likely to report experiencing discrimination as a student or faculty member (American Economic Association, 2019).

The Economics Job Market Rumors forum (EJMR), as its name suggests, was established to share information about job interviews and outcomes anonymously in each year's hiring cycle, though it is active year-round. According to a report by the forum administrator, about 80% of EJMR users who visit or post on the forum were males as of September 2017.[1] I scraped about 2.2 million posts from the first and the last page of each thread initiated or updated between October 2013 and October 2017 on this forum. Using a list of gender classifiers such as "she"/"he" from the most frequent 10,000 words in the EJMR postings, along with the names of over 9,000 active researchers and recent economics PhD graduates, I identify about 100,000 posts that discuss women (Female posts), and about 330,000 posts that discuss men (Male posts). About 63% of the threads in this four-year sample include at least one Female or Male post.

To guide my analysis, I develop a model of rumors that lays out a set of explanations for why people post differently about women and men in the profession. Posters are assumed to value their contribution to public knowledge about the relationship between professional characteristics and jobs in the profession. They are also assumed to gain utility by boosting the professional reputation of members of their own gender group relative to that of members of the opposite gender group. In any given thread, a poster can either reveal his or her private signal about the subject's professional ability or "change the subject" and discuss the subject's personal characteristics, adding noise to the discussion and clouding readers' assessments of the subject's true ability. As a result of the competing incentives, posters will tend to reveal

Received for publication June 25, 2018. Revision accepted for publication May 22, 2019. Editor: Shachar Kariv.
*Wu: Harvard University.
I am extremely grateful to David Card for his invaluable guidance throughout this project and to Janet Currie, Claudia Goldin and Amanda Pallais for many helpful discussions. I also thank Rediet Abebe, Patrick Button, Sydnee Caldwell, Krishna Dasaratha, Ellora Derenoncourt, Will Dobbie, Alessandra Fenizia, Nagore Iriberri, Wei Jiang, Patrick Kline, Michael Reich, David Romer, Jesse Rothstein, Dick Startz, Justin Wolfers, Esmée Zwiers, two anonymous referees, and seminar participants at UC Berkeley, Harvard Business School, Princeton, Bowdoin College, and AEA 2018 meeting for helpful comments. Bharvee Patel provided excellent research assistance. I gratefully acknowledge funding from the Lab for Economic Applications and Policy at Harvard. All errors are my own.
A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00877.

[1] The administrator of the EJMR forum released a statement in September 2017 claiming that 20% of EJMR users are female (https://www.econjobrumors.com/topic/kirk-statement-on-recent-events-and-moderation-policy). The number appeared to come from a third-party analysis of users' web-browsing cookies.

The Review of Economics and Statistics, December 2020, 102(5): 867–880

https://doi.org/10.1162/rest_a_00877

positive information about the abilities and accomplishments of members of their own affinity group and negative information about members of their out-group. Posters who receive a positive signal about a member of the out-group will instead post about personal characteristics, casting doubt on the professional accomplishments of the opposite group.

I test these ideas by measuring the differences in the topics of discussion in Female and Male posts and by quantifying the effects of gender on transitions between topics in the dynamics of a conversation. I begin by classifying the most frequent 10,000 words in the EJMR postings into different categories, grouping them into two broad topics: Academic/Professional and Personal/Physical. I record the number of words from each topic in each post and use the token count as a proxy for the extent to which the poster emphasizes the professional or personal characteristics of the subject. In August 2017, a *New York Times* article by Justin Wolfers raised concerns about the gendered content on EJMR and led to some changes in moderation efforts on the EJMR forum.[2] To take into account the influence of media coverage and moderation, I split the sample by whether a thread started before August 2017.

A direct comparison between gendered posts shows that Female posts on average contained 42% less Academic/Professional terms but 196% more Personal/Physical terms than Male posts prior to August 2017. I then break down these topic differences by job rank of the subject, which can be assigned to about 15% of gendered posts through keywords such as “job market candidate” or the name of the economist mentioned in a post. The gender gap in the number of terms related to professional characteristics is relatively smaller at the junior or senior faculty levels than at the graduate student and job market candidate levels but remains statistically significant. I also find that female posts continued to contain significantly more Personal/Physical terms at each job rank. These contrasts reveal a systematic tendency to deemphasize professional characteristics of women, which can be interpreted as a mechanism by which male posters boost their in-group identity and reinforce the perception of women as the out-group in the profession.

While the media coverage of EJMR in August 2017 led to some initial narrowing of the gender gap in the emphasis on professional characteristics, this pattern did not persist. Moreover, changes in EJMR’s moderation policies appear to have had little to no effect on the average number of Personal/Physical terms in Female versus Male posts. Both findings suggest strong inertia in stereotype beliefs about gender.

Since posters interact with each other within each thread, I present an empirical framework to measure gender bias in the dynamics of a conversation. I test whether a discussion about a female versus a male in a post systematically affects the likelihood that later posts focus on professional versus personal topics. Using a discrete choice model, I estimate the average marginal effects of gender on the probability of each possible transition between three states that represent the main topic of discussion in the thread: Purely Professional, Personal, or Others. I focus on posters who have decided to join an existing thread after seeing its most recent post and assume that the heterogeneity in posters’ preferences can be absorbed by observable characteristics of threads that posters select themselves into. Relative to the baseline where the prior post is not gendered (Genderless), I find a significant 2.3 percentage point difference-in-differences between the probability of deviating from a purely professional topic when the prior post is Female than when it is Male.

Although threads with more posts are less likely to get off track than those just started, a Female post still has a significantly higher chance of triggering a deviation from professional topics than a Male post does. Once such a deviation occurs, it is also significantly less likely to come back to professional topics from personal or other topics after a Female post. These gender differences in transition rates further support the hypothesis that an emphasis on a female subject’s professional characteristics can pose an identity threat to some male posters’ who will then mention nonprofessional attributes as a means to muddy the understanding about the subject’s true ability and protect their own identity in return.

Previous analysis of the EJMR forum documents the occurrences of explicitly sexual and discriminating terms associated with discussions about women that suggest an unwelcoming culture online (Wu, 2018). I add to this evidence by documenting a systematic tendency to deemphasize women’s professional accomplishments while highlighting their personal characteristics and by providing an identity-based interpretation of peers’ attitudes toward women in traditionally male-dominated fields that extend beyond this particular forum or the economics profession. The model of rumors in this paper is linked to the social identity theory in Tajfel and Turner (1986), which highlights a bias toward members of the insider group, and to the formal development of identity theory in economics in Akerlof and Kranton (2000) and more recently in Gennaioli and Tabellini (2019). The divergence in the portrayal of women and men along the professional versus personal dimensions is also consistent with the prediction of the model developed by Bordalo et al. (2016) in which stereotype bias leads to an exaggeration of the contrast between groups.

Taken as a whole, my findings suggest that at least some men are reluctant to let the public learn about the true distribution of women’s professional ability, which would be crucial to promote integration in a profession (Goldin, 2015). Finally, this paper is related to the literature on the link between gender role attitudes and women’s labor market outcomes (e.g., Fortin, 2005, 2015; Dahl, Kotsadam, & Rooth, 2018). The lack of progress in attitudes toward women as indicated by the EJMR forum can help explain part of the persistent gender gap in a profession.

[2] Wolfers (2017). The forum released an official statement regarding its new moderation policy in September 2017: https://www.econjobrumors.com/topic/kirk-statement-on-recent-events-and-moderation-policy.

II. A Model of Rumors

In this section, I develop a simplified model of rumors that lays out two incentives of posters to engage in anonymous discussions about other members in the profession, (a) contribution to the public knowledge and (b) identity boosting, and explains how the trade-off between these incentives can lead to stereotyping behavior that favors each poster's in-group while diminishing the out-group. I present only the key elements of this model and its main predictions. The details of the model are in appendix A.

Suppose in a profession there are two groups $\{F, M\}$, where the M group is more numerous and is traditionally considered the insider group. Given an anonymous message board and a thread of posts about a given subject (who is either F or M), a poster can either reveal his private signal about the subject's professional characteristics, which each poster derives from the subject's observable records such as publications and presentations, or discuss the subject's personal characteristics, which add uncertainty to the public's assessment of one's professional accomplishment. For example, under a discussion about a female economist's publications, a post that says, "She also had two kids during the 7 years hence the extended tenure clock," suggests that the subject's childbearing decision boosts her job promotion more than warranted by professional records alone and thus casts doubts on the subject's true ability.

I assume the utility of posting arises from two sources. First, a poster values his or her contribution to the public knowledge about the job market. The further is the poster's private signal from the opinion of other posters, the more utility he derives from moving the average perception closer to his own belief. In contrast, emphasizing personal characteristics is costly, as the poster makes the public information less precise. Second, following Akerlof and Kranton (2000), I incorporate a poster's identity relative to the subject into the utility function. Specifically, a poster perceives an identity threat if the subject from the out-group has higher professional characteristics than his own but affirms a positive image of himself if the subject comes from his in-group. These assumptions are consistent with a key argument in social identity theory that people aim to achieve a positive image of their own group in contrast with the opposite group (Tajfel & Turner, 1986).

This framework yields a set of predictions about posting behavior toward different groups:

1. A poster who cares about identity tends to reveal positive signals about the professional characteristics of subjects from the in-group but hides positive signals about subjects from the out-group. Revealing a positive professional signal about a subject from the same group enables a poster to contribute to the public knowledge and simultaneously perceive a more promising professional identity of his own. In contrast, when the subject comes from the out-group, he faces a trade-off between enhancing the public understanding about the job market and protecting his self-image relative to the out-group. Appendix figure A1 illustrates the contrast between the group-specific ranges of private signals over which a poster is willing to reveal. Without concerns about identity, each poster will be equally likely to reveal professional signals about subjects from in- versus out- groups.
2. A poster who cares about identity tends to discuss personal characteristics of subjects from the out-group but not from the in-group. Discussions about the subject's personal characteristics increase the uncertainty about his or her true professional ability and thus are costly to posters who are assumed to care about their contributions to knowledge about the job market. However, by avoiding admitting to a positive professional portrayal of someone from the out-group, a poster can protect his identity in comparison. This trade-off leads to a higher emphasis of personal characteristics of the out-group rather than the in-group.
3. The more a poster cares about identity, the larger the gap between the average professional signals he reveals about subjects from in-group versus out-group and the more likely he discusses personal characteristics of the out-group than the in-group. Appendix figure A2 illustrates that when a poster puts more weights on the self-image, there is a higher divergence in the average professional signals he reveals between the two groups, and it follows that the gap in the average personal signals revealed is also larger.
4. When posters take into account how others would react to their remarks, those selected into posting either hold very different views from other posters or are relatively more sensitive to their identity relative to the subject. Posters are discouraged from expressing outrageous opinions when they are concerned about an immediate backlash against them. I use a simple extensive game as in Akerlof and Kranton (2000) to illustrate this point (figure A3). Combined with the third prediction, this result suggests that the gaps in the professional and personal signals revealed on the forum are exaggerated by posters who hold stronger views and are more vulnerable to identity threats from the out-group.

I test for the first two predictions (corresponding to proposition 1 in appendix A) from both a static and a dynamic perspective. Assuming that the majority of EJMR posters are male, if the predictions were true, the data would show a higher emphasis on professional characteristics of men than women and a higher emphasis on personal characteristics of women than men. Since each thread environment is dynamic and interactive in nature, a positive signal about a woman's professional ability in one post is more likely to trigger a transition toward her personal characteristics in future posts than that of a man's professional ability.

FIGURE 1.—NUMBER OF NEW THREADS PER MONTH

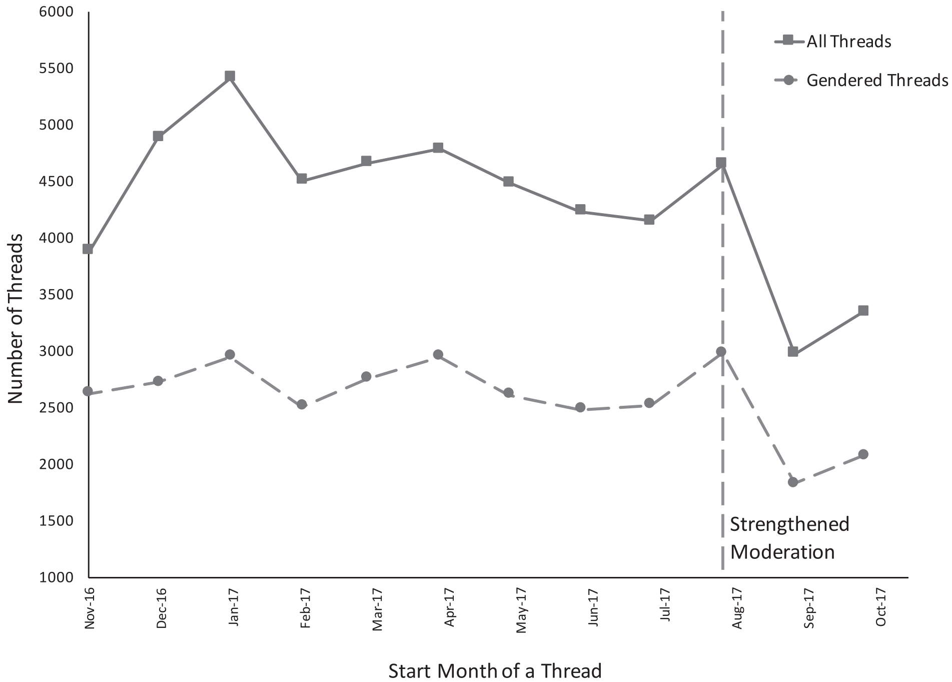

This figure shows the number of new threads initiated in each month between November 2016 and October 2017, in the full sample and the gender sample (threads that include at least one Female or Male post), respectively. For threads started before November 2016, I cannot identify the calendar month from the rough time stamps such as "1 year ago," "2 years ago" listed on EJMR. In August 2017, a *New York Times* article by Justin Wolfers raised concerns about the gendered content on the EJMR forum and led to strengthened moderation policies by EJMR, which appeared to result in the removal of a significant number of threads in September and October 2017.

The selection of posters is not testable in the data, but it provides an explanation for the prevalence of stereotyping behavior that exaggerates the true differences between women and men in the profession. Anonymity presumably aggravates the selection as it enables posters with more biases to voice their opinions without the constraints of social pressure as in other public settings. As a result, the professional information about women is systematically negatively biased relative to that of men on the forum, and the discussions about personal characteristics make the public less certain about the subject or the entire group's professional abilities, which slows the information updating about the underrepresented group that is crucial to the integration in the profession under the pollution theory in Goldin (2015).

III. EJMR Data

As of October 28, 2017, there were 306,253 threads on the EJMR forum originating over the previous seven years. The threads are organized in reverse chronological order, by the time of each thread's latest post. Figure 1 shows that the number of new threads per month peak in December and January when candidates finish academic job applications and employers start to arrange interviews and fly-outs, but the forum remains active in other months.

I took two steps to create my data set. First, I scraped the main pages of the forum. At the time of my data extraction, there were 8,759 pages. A typical page contains 35 threads, and it records each thread's title, the time of the latest update, the number of posts, the number of views, and the votes by users (see appendix figure B1). I then scraped the posts on the first and last page of each thread initiated or updated between October 2013 and October 2017.[3] In this way, I obtained a data set of 2,217,046 posts (including titles) across 223,475 threads.

In the absence of a preexisting dictionary, I identified the most frequent 10,000 words from the raw text and recorded the word counts for each word in each post. Based on this list I constructed measures of topics—{Academic/Professional, Personal/Physical}—gender, and, if possible, job rank—{graduate student, job market candidate/postdoc, junior faculty, senior faculty} of the subject of discussion (see appendix B).

A. *Identifying Gendered Posts*

Among the most frequent 10,000 words, 53 indicate a post about a female (e.g., "she"/"woman") and 204 words that indicate a post about a male (e.g., "he"/"man"). Appendix table B1 provides a complete list of such gender classifiers. The imbalance in the number of classifiers arises from the

[3] A typical thread contains at most twenty posts on each page (see appendix figure B2 for a screenshot of a thread). Among threads in the past four years, about 12% exceeded one page, and about 4% exceeded two pages at the time of data extraction.

different numbers of gendered first names or economists' last names among the top 10,000 words. Using these classifiers, I identified 102,956 posts that discuss women (Female posts) and 327,670 posts that discuss men (Male posts). About 10% of Female/Male posts include classifiers of both genders, and they have been reclassified through a Lasso-regularized logistic model that predicts gender through counts of the most frequent 10,000 words excluding the gender classifiers.[4]

To address the imbalance in the number of economists' names among the 10,000 words and further identify the posts about specific economists, I assembled a list of 5,003 authors of National Bureau of Economic Research (NBER) working papers from 2014 to 2017 and a list of 4,724 job market candidates who graduated from 36 top economics PhD programs in the United States and Canada from 2011 to 2018. Table B2 summarizes the number of female and male economists in each sample.

The sample of NBER authors comprises active researchers in the economics profession. Since 2014, the administrators of the EJMR forums post abstracts of new working papers from NBER every week. I scraped information about 5,003 authors of 4,478 working papers from the NBER website, among whom 1,008 are affiliated with NBER as research associates, 301 are faculty research fellows, and the rest are their collaborators.[5]

Junior economists are less likely to be affiliated with NBER and are thus underrepresented in the NBER sample. To address this selection issue, I collected an additional sample of recent job market candidates from top economics programs between 2011 and 2018. I focused on institutions in the United States and Canada that ranked among the top fifty economics departments by econphd.net in 2004.[6] I found 4,724 PhD graduates from these institutions on the ProQuest database of doctoral dissertations, and lists of job market candidates or placement records from the department websites. Table B3 provides a list of schools in this sample.

To identify the gender of each person, I first matched his or her full name with the data set of 48,000 economists with gender assignment assembled by Card et al. (2019). For those who were not matched, I used the "gender" and "genderizeR" packages in R to predict gender from first names and assign gender only if the predicted probability of being a female or male was at least 0.85. Finally, manual searches and assignments were done on the remaining 1,200 economists.

I searched each person's full name in the sample of 2.2 million EJMR postings. If a post includes one's full name, I then searched her first name, last name, and initials within the same thread of this post and therefore identified more posts that discussed this person. In this way, I found 57,816 posts that mentioned NBER authors and 16,739 that mentioned job market candidates. About 87% of posts that include economists' names have already been picked up by gender classifiers, as these posts are likely to include pronouns like "he" or "she." An advantage of identifying gender through names, however, is that I can collect information about the job rank of the subject of discussion and explore patterns across the job ladder in later sections.

Table 1.—Sample Overview

	Female	Male
Before August 2017		
Number of posts	99,659	318,873
Number of threads	41,243	116,996
Number of posts by job rank		
Graduate students	3,111	11,359
Job market candidates/postdocs	2,156	8,932
Junior faculty	2,335	9,675
Senior faculty	2,097	18,339
August–October 2017		
Number of posts	4,817	15,848
Number of threads	2,129	6,192
Number of posts by job rank		
Graduate students	181	639
Job market candidates/postdocs	112	374
Junior faculty	159	718
Senior faculty	131	1,059

Panel A reports posts under threads initiated before August 2017, whereas panel B reports posts under threads initiated between August and October 2017. "Number of Threads" records the number of threads that contain at least one Female or Male post, respectively.

In summary, I found 104,476 Female posts and 334,721 Male posts in total, comprising over 20% of all posts during the sample period. These gendered posts come from 139,981 threads, representing about 63% of all threads in the sample. Table 1 further breaks down the sample by identified job rank, before and after August 2017, when there was a change in the forum's moderation policy.

IV. Topic Differences between Gendered Posts

To measure the topics at the post level, I manually classified the most frequent 10,000 words into fifteen categories. Table B4 explains how I grouped certain categories to consider two main topics of interest: (a) Academic/Professional and (b) Personal/Physical. The first topic is consistent with the professional purposes of the EJMR forum, whereas the second topic emphasizes personal characteristics, which add noise to professional discussions and to some extent reflect posters' stereotyping behavior under the model of rumors.

For each post, I count the number of occurrences of words from each topic, which represents the degree to which a post emphasizes a given type of characteristic of the subject. Table 2 displays the differences between Female and Male posts in the mean number of Academic/Professional words, the mean number of Personal/Physical words, the fraction of posts related to each topic separately, and finally the fraction of posts

[4] I trained a Lasso-logistic model with fivefold cross-validation on 75% of posts that refer uniquely to one gender or the other, and then selected the optimal p-score threshold that minimizes the mean squared error for predicting gender on the remaining 25% as a test set. The model and the training process are discussed in detail in Wu (2018).

[5] As mentioned on NBER's website (https://www.nber.org/info.html), research associates are tenured faculty at their home institutions, and their appointments at NBER are approved by the NBER board of directors, whereas faculty research fellows are typically junior faculty.

[6] The ranking of economics departments can be found at http://econphd.econwiki.com/rank/rallec.htm.

TABLE 2.—TOPIC DIFFERENCE BETWEEN FEMALE AND MALE POSTS

	Female	Male	Difference	SE
A. Before August 2017				
Counts				
Mean number of Academic/Professional terms	1.8792	3.2468	−1.3676	(0.0236)
Mean number of Personal/Physical terms	0.9996	0.3379	0.6617	(0.0085)
Indicators				
Has any Academic/Professional term	0.4772	0.5925	−0.1153	(0.0023)
Has any Personal/Physical term	0.4396	0.1949	0.2446	(0.0022)
Purely Academic/Professional	0.2387	0.4595	−0.2208	(0.0020)
B. August–October 2017				
Counts				
Mean number of Academic/Professional terms	2.6425	3.525	−0.8825	(0.1274)
Mean number of Personal/Physical terms	0.8943	0.3423	0.552	(0.0342)
Indicators				
Has any Academic/Professional term	0.5553	0.5968	−0.0415	(0.0102)
Has any Personal/Physical term	0.3971	0.1904	0.2068	(0.0096)
Purely Academic/Professional	0.3104	0.4597	−0.1494	(0.0097)

This table shows the topic differences between Female and Male posts, measured by counts of words in each topic and indicators for containing any word from a given topic. Standard errors in the last column are robust and clustered at the thread level. Panel A reports posts under threads initiated before August 2017, whereas panel B reports posts under threads initiated between August and October 2017.

FIGURE 2.—TRENDS OF TOPICS IN FEMALE VERSUS MALE POSTS

(a) Mean Number of *Academic/Professional* by the start month of each thread

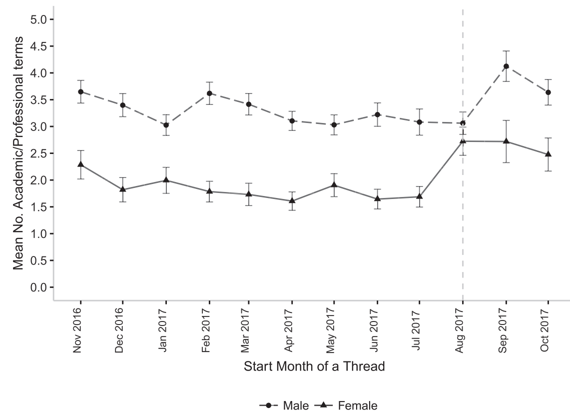

(b) Mean Number of *Personal/Physical* by the start month of each thread

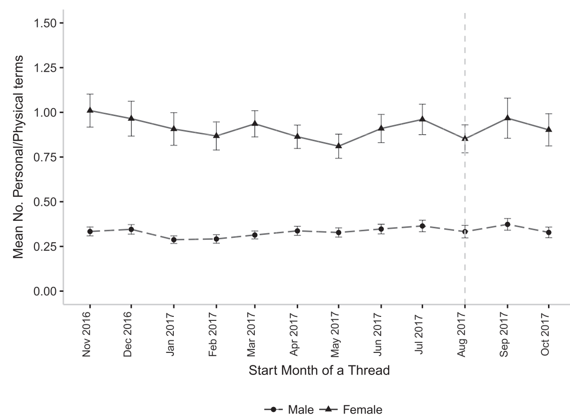

This figure plots the sample means (95% CI shown) of the number of Academic/Professional or Personal/Physical terms in Female versus Male posts from threads initiated within each month between November 2016 and October 2017. I identify the start month of each thread by the rough time stamp of its first post. The dashed line at August 2017 indicates the beginning of media coverage and strengthened moderation policies on the EJMR forum.

that are purely Academic/Professional.[7] The standard error for each measure of topic difference is clustered at the thread level.

Panel A in table 2 shows that prior to August 2017, on average there were 3.25 Academic/Professional terms in Male posts but 1.37 significantly fewer such terms in Female posts. Figure 2a further breaks down this gap by the month in which a thread was started, identified among threads initiated between November 2016 and October 2017.[8] The gender gap in terms of a percentage difference fluctuated between 34% and 51% before August 2017, and it did not show significant differences between the job market season and other months.

The other topic, Personal/Physical, gives a different picture. On average, Female posts contained about one word concerning personal information or physical appearance, more than double the means across Male posts. Although the magnitude of this difference seems smaller than that in the number of Academic/Professional words, it is worth noting that this category includes a significant fraction of words related to physical attributes or sexual content that objectify women and reinforce the perception of them as an out-group in the profession.

These topic differences show that the overall population of posters put a significantly lower emphasis on women's professional characteristics than on men's and a significantly higher emphasis on women's personal characteristics. Both patterns are consistent with the predictions in the model of rumors (section II) from a static perspective. I interpret discussions about women's personal characteristics as a means to cast doubt on their professional abilities and thus protect male posters against an identity threat.

A. *The Role of Moderation Policies*

I split the data by whether a thread was initiated before August 2017, when a *New York Times* article by Justin Wolfers drew attention to gender issues in the economics profession and triggered a strengthening of EJMR's moderation policies that removed controversial or inappropriate content on the forum. Figure 1 shows that the total number of new threads first increased in August but then dropped by about 36% in September when the forum put a new moderation policy into effect. In particular, there were 38% fewer new threads

[7] A post is considered purely Academic/Professional if it contains at least one term from the professional topic but none from the personal topic.

[8] I used the rough time stamp under the first post of each thread to identify the month in which a thread was initiated. The time stamps are written as "1 day ago," "1 month ago," "11 months ago," "1 year ago," "2 years ago," and so on. Therefore, I was only able to identify the start month of threads between November 2016 and October 2017, within one year as of my latest web scraping.

related to gender in September than in August. This pattern suggests that the content since August 2017 was more selective than before and might not be representative of the views of the original population of posters.

Panel B in table 2 shows that between August and October 2017, there was a noticeable shrinkage of the gender gap in the Academic/Professional topic, which can be attributed to both a rise in the emphasis on professional attributes in Female posts and a selection of threads due to new moderation. Figure 2a provides a more nuanced picture: during August 2017, Female posts contained 2.7 academic terms on average, 11% less than the mean across Male posts. However, the gender gap exceeded 30% again in the next two months.

Figure 2b shows that the gender difference in the Personal/Physical topic remained significant in each month between August and October 2017. The changes in EJMR's moderation policies since August 2017 did not make an immediate change on the Personal/Physical topic as it did for the Academic/Professional topic, suggesting that it is particularly difficult to break the association between women and nonprofessional and stereotypical discussions.

FIGURE 3.—GENDER DIFFERENCES IN ACADEMIC/PROFESSIONAL BY JOB RANK

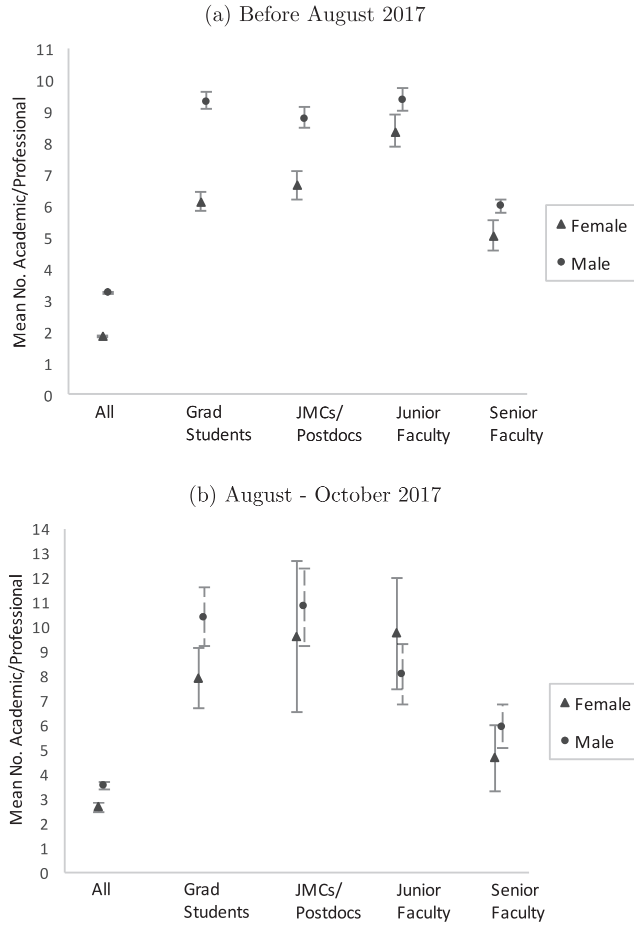

This figure plots the sample means (95% CI shown) of the number of Academic/Professional terms among all Female versus Male posts and among Female versus Male posts at each job rank, assigned to about 15% of all posts in the sample (see appendix B).

B. *Gender Gaps by Job Rank*

To examine whether these topic differences vary across positions on the job ladder, I use information about specific economists and a list of keywords to identify the job rank in each post (see appendix B for details). I focus on four observable ranks: Graduate Students, Job Market Candidates and Postdocs, Junior Faculty, and Senior Faculty. About 15% of all gendered posts are assigned a job rank. Table 1 summarizes the number of Female and Male posts at each job rank.

Figure 3a displays the mean number of Academic/Professional terms at each job rank among posts before August 2017. In comparison with the full sample, both Female and Male posts with assigned job ranks contained more Academic/Professional terms on average. However, the gender gap at each rank was significant before August 2017. For example, a typical post about female job market candidates or postdocs had about 6.66 Academic/Professional words, 2.14 fewer ($t = -7.63$) than a typical post about male candidates. The gap shrank in both absolute and relative terms for junior and senior faculty but remained statistically significant. The change in EJMR's moderation policies and other factors around the media exposure in August 2017 appeared to reduce this gender gap, as shown in figure 3b, most strikingly among discussions about job market candidates and junior faculty.

Figure 4 shows that the emphasis on personal characteristics remained significantly higher in Female posts than in Male posts across all job ranks, and strengthened moderation policies did not make a notable difference for this topic. Posts about senior faculty of each gender contained fewer words concerning personal information or physical appearance. Assuming the majority of EJMR posters are early in their career, an identity-based interpretation of this pattern is that posters put more weight on their self-image relative to subjects closer to them on the job ladder, and thus they are less likely to feel professionally threatened by senior economists. In contrast, the gap became significant again among junior faculty, many of whom were evaluated by EJMR posters over whether they deserved tenure. The emphasis on a female assistant professor's personal characteristics can be interpreted as adding noise to the public assessment of her professional ability, which potentially helps a poster maintain a relatively higher professional status of his own.

V. Dynamics of Topics in Sequential Conversation

Moving beyond the static analysis of topic differences between gendered posts, I present an empirical framework to measure stereotyping in the dynamics of a conversation. Each thread on the forum is a dynamic environment in nature where posters interact with each other. In the model of rumors, a

FIGURE 4.—GENDER DIFFERENCES IN PERSONAL/PHYSICAL BY JOB RANK

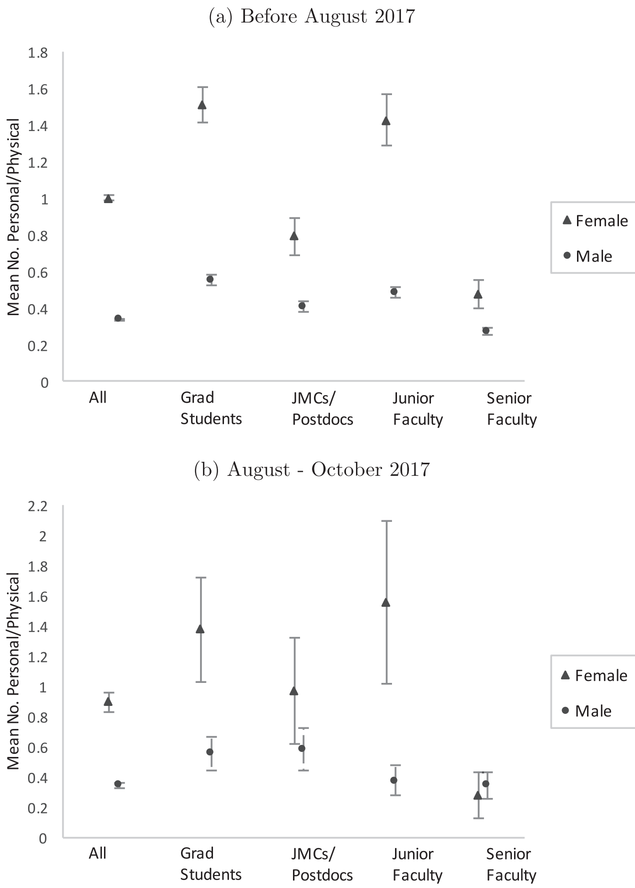

This figure plots the sample means (95% CI shown) of the number of Personal/Physical terms among all Female versus Male posts, and among Female versus Male posts at each job rank, assigned to about 15% of all posts in the sample (see appendix B).

poster's choice of topic depends not only on his or her private information about the subject, but also on the signals revealed by previous posters and the expected reactions from future posters (see appendix A). A model without consideration of the interactions between posters would not be able to capture the autocorrelation of topics within a thread and how they vary once attention turns to one gender or the other.

I use a discrete choice model to estimate the effects of gender on the transition probabilities between topics in gendered threads that contain at least one Female or Male post. I interpret the transitions as a means to contribute to the public understanding about the job market and boost a poster's professional identity relative to the subject. Since about 80% of EJMR users (visitors and posters) were claimed to be male as of September 2017, I consider the patterns discussed below to be primarily driven by the preferences of male posters.[9] Table 3 provides some examples of consecutive posts in actual EJMR threads that illustrate the transitions between topics when the prior post is gendered.

A. *Choice between Discrete States of Posts*

I focus on the decision making of a poster regarding the topic of a new post added to an existing thread t with N number of posts. There are three possible states of each post, $s(n) \in$ {Purely Professional, Personal, Others}, which lead to nine possible transitions between the states of consecutive posts. Each post can be gendered (Female or Male), or not gendered (Genderless).

Let $\mathbf{Gender_{t,N}}$ denote a vector of indicators for whether the last post observed by the poster discusses a female or a male, respectively. Conditional on the Nth post of thread t in state s_0, I specify the utility function of poster k creating the $(N+1)$th post in state s as

$$u_{ks} = \alpha_{ks} + \beta_{ks}\mathbf{Gender_{t,N}} + \epsilon_{ks}, \tag{1}$$

where α_{ks} represents k's utility from writing a new post in state s in reaction to a Genderless post, and the vector β_{ks} captures the additional utility k obtains when the Nth post is Female or Male, respectively. I assume the error ϵ_{ks} is independently and identically distributed according to type 1 extreme value across posters and choices.

In this anonymous setting, I do not have information about individual posters, but I proceed by assuming that posters who select themselves into the same type of threads have the same preferences for topics. Following the modeling assumptions in Bayer, Ferreira, and McMillan (2007), I allow each poster's preferences for topics to vary with a set of observable thread characteristics denoted by $\mathbf{Z_{t,N}}$, according to

$$\alpha_{ks} = \alpha_s + \pi_\alpha \mathbf{Z_{t,N}},$$
$$\beta_{ks} = \beta_s + \pi_\beta \mathbf{Z_{t,N}},$$

where (α_s, β_s) is shared by all posters and (π_α, π_β) capture the heterogeneity in preferences under different types of threads. The variables in $\mathbf{Z_{t,N}}$ control for initial conditions that indicate the topic (Purely Professional, Personal or Others) and the gender of the subject (Female, Male, or Genderless) in the title and the first post of the thread. Initial conditions should be taken into account if they are assumed to be correlated with any unobserved permanence in a dynamic decision process (see Eckstein & Wolpin, 1989; Keane & Wolpin, 1997; see also Aguirregabiria & Mira, 2008, for a survey). In particular, the initial state is considered important in shaping the theme and triggering subsequent discussion in recent studies about the dynamics of online conversation (see Farajtabar et al., 2015, 2017; Rizoiu et al., 2017).

[9]The administrator of the EJMR forum released a statement in September 2017 claiming that 20% of EJMR users are female (https://www.econjobrumors.com/topic/kirk-statement-on-recent-events-and-moderation-policy). The number appeared to come from a third-party analysis of users' web-browsing cookies. Note that "users" includes all visitors of the forum. It is not clear whether female and male users have the same propensity to post on the forum.

TABLE 3.—EXAMPLES OF TRANSITIONS BETWEEN TOPICS IN CONSECUTIVE POSTS

*N*th Post		(*N* + 1)th Post		
Gender$_N$ $s(N)$	Content	*Gender*$_{N+1}$	$s(N+1)$	Content
From Purely Professional				
Male	"I think [Name] is the best. **He** has the most solid job market paper in IO [industrial organization] …"	*Genderless*	Purely Professional	"Agreed. Will certainly be best IO candidate on market"
Female	"This is a very weak record especially given the fact that **she** took 9 years to get tenure."	*Female*	Personal	"Collegial externalities—**she** looks nice, great gender."
Female	"**Her** quantity is pretty outstanding. 3 publications, 4 working papers …"	*Genderless*	Others	"Stop with the self-promotion you little shiets"
From Personal				
Male	"All that matters for **men** is what shows in a dress shirt … hygiene shows best."	*Genderless*	Purely Professional	"It's really all about the JMP [job market paper]."
Female	"When I asked them why that teacher got good evaluations the student literally said it was because: '**she** was hot!'"	*Genderless*	Personal	"Just put on a short dress, maybe show a little panties and at the very least show a panty line and then act cute and you'll do fine."

This table provides examples of transitions of topics between consecutive posts (from the *N*th post to the $(N+1)$th post) in actual threads from EJMR. *Gender* represents the gender of the subject in a post (Genderless if neither Female nor Male), *s* refers to the state or main topic of a post (Purely Professional, Personal, or Others), and the content is abbreviated for illustrative purposes. For Female or Male posts, the gender classifiers (female or male) for each post contains are in specified.

To test for posters' preferences over discussions about specific job ranks, $\mathbf{Z_{t,N}}$ includes an indicator for each job rank along with a group without rank assigned in the *N*th post. Preferences over length of existing threads are captured by $\ln(N)$—the log number of previous posts, which allows for differential returns to posters when switching topics in shorter versus longer threads. Finally, $\mathbf{Z_{t,N}}$ includes the fraction of posts under each possible state in thread t, which do not vary across posts and presumably absorb any remaining unobserved thread-level propensity for each possible transition of topics.[10]

Given the specification above, I rewrite equation (1) as

$$u_{ks} = (\alpha_s + \pi_\alpha \mathbf{Z_{t,N}}) + (\beta_s + \pi_\beta \mathbf{Z_{t,N}})\mathbf{Gender_{t,N}} + \epsilon_{ks}, \tag{2}$$

Under the assumption that ϵ_{ks} is distributed type 1 extreme value, the problem can then be estimated as a multinomial logit. When each poster chooses the state that maximizes equation (2), the realized choice probabilities are

$$P(s(N+1) = s | s(N) = s_0) = \frac{\exp(u_{ks})}{\sum_{s'} \exp(u_{ks'})}.$$

I begin by estimating the average marginal effects of gender in the prior post on the probability of transitioning from state s_0 to state s in the current post. It is particularly interesting to examine the gender differences in the persistence in professional topics and the switches between professional and personal topics. The incentive to affirm a poster's own status may lead him to continue emphasizing the professional characteristics of subjects similar to him, whereas the incentive to improve one's self-image relative to subjects from the opposite group is likely to result in a deviation from the professional topics (see table 3 for examples). In section VC, I also discuss some alternative explanations for my findings.

B. Main Results

I estimate the model on 132,936 gender-related threads initiated or updated before August 2017 to avoid contaminating the results by heavily censored content since then.[11] The sample includes 99,659 Female posts, 318,873 Male posts, and about 1.1 million Genderless posts in total.[12]

Figure 5 shows the average marginal effects of gender on the probability of each possible transition, conditional on the state of the previous post. Standard errors are clustered at the thread level to take into account the correlation between posts within the same thread. All the estimates are relative to the base group comprising transitions from posts that are not gendered (Genderless).

First, there is a relatively lower persistence in purely professional topics when the prior post mentions a female, but this is not the case following the mention of a male. Conditional on the prior post being purely professional, the current post is on average 2.9 percentage points ($t = -6.5$) less likely to stay on the professional topic when the prior one is Female rather than Genderless, but 1.0 percentage point ($t = 3.9$) more likely to switch to a personal topic and 1.9 percentage points ($t = 4.4$) more likely to switch to miscellaneous topics not identified as professional or personal. Yet

[10] As discussed in section IVB, about 15% of gendered posts are assigned one of the four job ranks based on reference to a specific economist or keywords: Graduate Students, Job Market Candidates/Postdocs, Junior Faculty, and Senior Faculty. The controls include an indicator for the 85% of the sample without job rank assigned, and this group is used as base.

[11] By "gender-related," I mean each thread contains at least one Female or Male post.

[12] Since I only scraped posts on the first and the last pages of each thread, I drop the "transition" from the last post on the first page and the first post on the last page if there are missing pages in between. Part of the difference between posts on the first page versus those on the last page is absorbed by the control $\ln(N)$ where N is the total number of previous posts, including those I did not scrape.

FIGURE 5.—AVERAGE MARGINAL EFFECTS OF GENDER ON ALL POSSIBLE TRANSITIONS

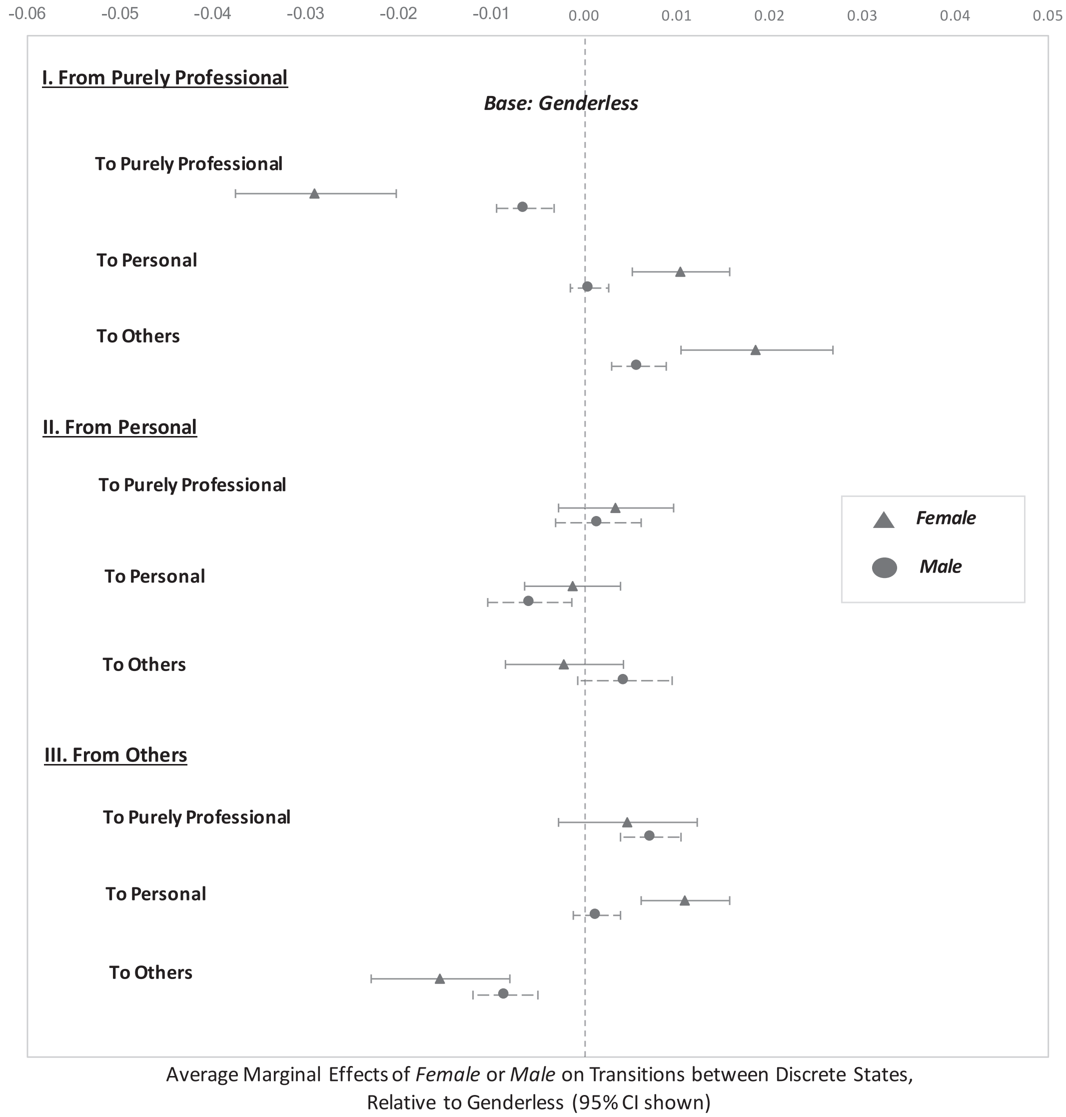

I estimate a multinomial logit model for transitions in topics on all threads started before August 2017 with at least one gendered (Female or Male) post. Standard errors are robust and clustered at the thread level.

when the prior post is Male, the deviation from a purely professional topic relative to the base group is 0.6 percentage points ($t = -4.1$), a much smaller difference compared with the Female group.

Figure 6 further reveals that the gender gaps in transitions from professional topics are salient at all different lengths of threads. When a thread just begins with one or two posts, a new post is about 4 percentage points significantly less likely to be persistent in professional topics following the mention of a female and about 2 percentage points significantly less likely so following the mention of a male (see figure 6a). As the thread gets longer, the probability of leaving professional topics shrinks for both genders, suggesting a decline in the leverage of new posters over a thread that has been going on for awhile. Nevertheless, even in threads at the 90th percentile with about 23 existing posts, the relative likelihood to deviate from purely professional topics after a Female post does not become insignificant, as it does after a Male post. Figure 6b shows that about 40% of the deviations from professional topics go to personal ones after a Female post, and such a tendency remains significant at the 5% level until the 90th percentile. Male posts, in contrast, do not trigger a significant transition toward personal topics at any length of thread on average.

Once a deviation from the professional topic occurs, it is less likely to come back to the professional track in discussions about women than those about men. Figure 5 shows that among transitions from personal topics, there is on

FIGURE 6.—AVERAGE MARGINAL EFFECTS OF GENDER ON TRANSITIONS BY LENGTH OF THREADS

(a) Staying on Purely Professional Topics

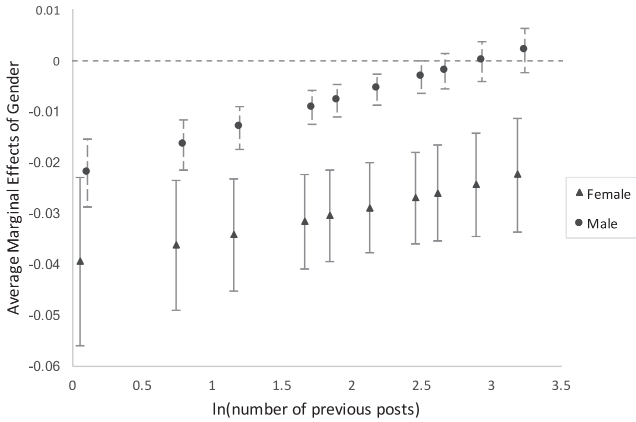

(b) Moving from Purely Professional to Personal Topics

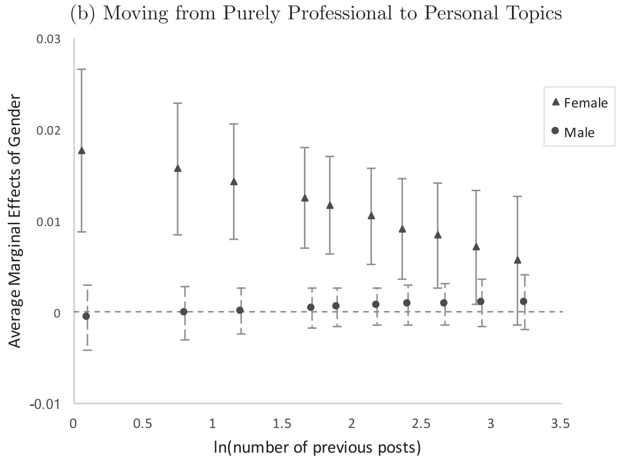

This figure shows the average marginal effects (with 95% confidence intervals) of the mention of a female or a male on the probability of (a) staying on the purely professional topics and (b) moving from professional to personal topics, relative to the baseline where the prior post is not gendered, at different lengths of existing threads. The minimum of the log number of previous posts at 0 represents the second post of each thread, at which the first possible transition of topics between consecutive posts occurs in a thread. The other data points are located at the deciles of the log number of previous posts (denoted by $\ln(N)$ in text), from 10% to 90%, respectively. Standard errors are robust and clustered at the thread level.

average a 0.6 percentage point significantly higher chance of escape from personal to purely professional or other topics when the prior post is Male rather than Genderless, while the Female group does not show any significant difference from the baseline. The Male group also shows a 0.7 percentage point significantly higher probability of moving back to professional topics from miscellaneous ones, in contrast with a 1.1 percentage point significant increase in the probability of switching to personal topics instead when the prior post mentions a female.

As the majority of posters are male, the lack of persistence in professional topics following a Female post is consistent with the main prediction from the model of rumors: male posters are less likely to reveal signals about a female subject's professional characteristics when their own identity is threatened by the out-group. The transition from professional to personal topics is arguably inappropriate in a professional setting and to some extent represents a belief that personal characteristics provide a premium to women's career (see the second example in table 3). Finally, the difficulty in moving back to professional topics from nonprofessional ones after a Female post suggests that stereotypical views can be easily reinforced on the forum, leading to a systematic deemphasis of professional characteristics of women as the out-group in the profession.

C. *Alternative Explanations*

I test for two alternative hypotheses that may explain the divergence in the transition rates between topics.

Selection of posters into different threads. Posters who select themselves into threads that start off with a personal topic related to women are unlikely to move the discussion toward a professional topic, whereas those who join a thread about a man's professional attributes are more likely to stay on the initial professional topic. The heterogeneity in posters' preferences over threads can contribute to the gender gaps in transitions in the data. In the discrete choice model, I assume posters' preferences are absorbed by the controls for thread characteristics, including initial topic and gender and mean topics across all posts. In particular, posters are prompted to click on a thread by its title listed on the main sites of the EJMR forum (see appendix figure B1). A typical title contains fewer than ten words, but it is arguably sufficient in conveying whether the thread means to be professional or personal and whether it pertains to a particular gender.

Table 4 reports the estimated average marginal effects of gender on transition probabilities from purely professional topics under different types of thread titles. In threads with purely professional titles, the mention of a female in the middle of a discussion leads to a significant 2.8 to 4.2 percentage point decrease in the probability of staying on purely professional topics, whereas the mention of a male shows insignificant difference from the baseline where the prior post is not gendered. The gender gap in this dimension is most salient in threads that are initially professional and mention a male in the title. Posters selected into male-oriented threads may be particularly reluctant to see professional discussions about women, which are perceived as a "pollution" to their own group and pose an identity threat. Nevertheless, the selection hypothesis does not explain why there is also a significant lack of persistence in professional topics under initially gender-neutral threads.

Under threads that begin with a personal topic and a female subject, a purely professional Female post is 3.0 percentage points more likely to trigger a transition toward personal topics; however, this transition rate remains significantly positive under initially professional threads. That is, the tendency to deviate from professional to personal topics is more systematic than what can be explained by selection of posters into different threads.

TABLE 4.—AVERAGE MARGINAL EFFECTS OF GENDER ON TRANSITIONS UNDER DIFFERENT INITIAL CONDITIONS

	(1) Staying on Purely Professional		(2) Professional → Personal	
	Female	Male	Female	Male
By characteristics of titles				
Purely Professional and Genderless	−0.0288	−0.0067	0.0106	0.0013
	(0.0051)	(0.0020)	(0.0033)	(0.0014)
Purely Professional and Female	−0.0277	−0.0105	0.0274	−0.0007
	(0.0103)	(0.0110)	(0.0059)	(0.0057)
Purely Professional and Male	−0.0415	0.0014	0.0153	0.0008
	(0.0125)	(0.0040)	(0.0074)	(0.0026)
Personal and Genderless	−0.0226	−0.0089	0.0152	−0.0001
	(0.0122)	(0.0074)	(0.0058)	(0.0036)
Personal and Female	−0.0203	−0.0130	0.0304	−0.0018
	(0.0152)	(0.0128)	(0.0076)	(0.0057)
Personal and Male	−0.0351	−0.0006	0.0195	−0.0005
	(0.0171)	(0.0083)	(0.0086)	(0.0041)

This table displays the average marginal effects of the prior mention of a female or a male on the probability of (a) staying on purely professional topics and (b) moving from purely professional to personal topics, relative to the baseline where the prior post is not gendered. Standard errors in parentheses are robust and clustered at the thread level. The initial conditions listed here are solely determined by the topic and gender in the title of each thread, which posters can see before clicking on the thread. The full set of initial thread characteristics in the model also include the topic (Purely Professional, Personal, or Others) and gender (Female, Male, or Genderless) in the first post of each thread.

In summary, the heterogeneity in the estimates suggests that the selection of posters into different types of threads does play a role in driving the gender gaps in transitions, but there are systematic patterns across all types of threads that cannot be explained by selection alone.

Valuations for knowledge about different job ranks. About 15% of gendered posts are assigned a job rank through information about economists they discuss or keywords such as "job market candidates" and "junior faculty" (see appendix table B5). Posts that discuss a specific job rank tend to include more Academic/Professional terms on average than the rest of the sample (see figure 3). If posters value the public knowledge about the profession and in particular the job market for both genders, there should be fewer gender differences in deviations from purely professional topics.

Figure 7 shows the average marginal effects of gender on transitions evaluated at four job ranks: Graduate Students, Job Market Candidates/Postdocs, Junior Faculty, and Senior Faculty. At each job rank, a Male post does not trigger a significant deviation from purely professional topics relative to the baseline (Genderless), whereas a Female post leads to a 3 to 5 percentage point significant decrease in persistence in professional topics. There is a significant gender difference in persistence at the junior faculty level, where the effect of a Female post is most negative. However, a null hypothesis of no gender difference in persistence cannot be rejected at other ranks, partially because the estimates for the Female group are noisier because there are fewer Female posts identified as each rank than there are Male posts (see table 1). The estimates for the effects of Female posts on transitions from professional to personal topics are also imprecise, and thus I cannot reject a null hypothesis of no gender difference at the 5% level.

FIGURE 7.—AVERAGE MARGINAL EFFECTS OF GENDER ON TRANSITIONS BY JOB RANK

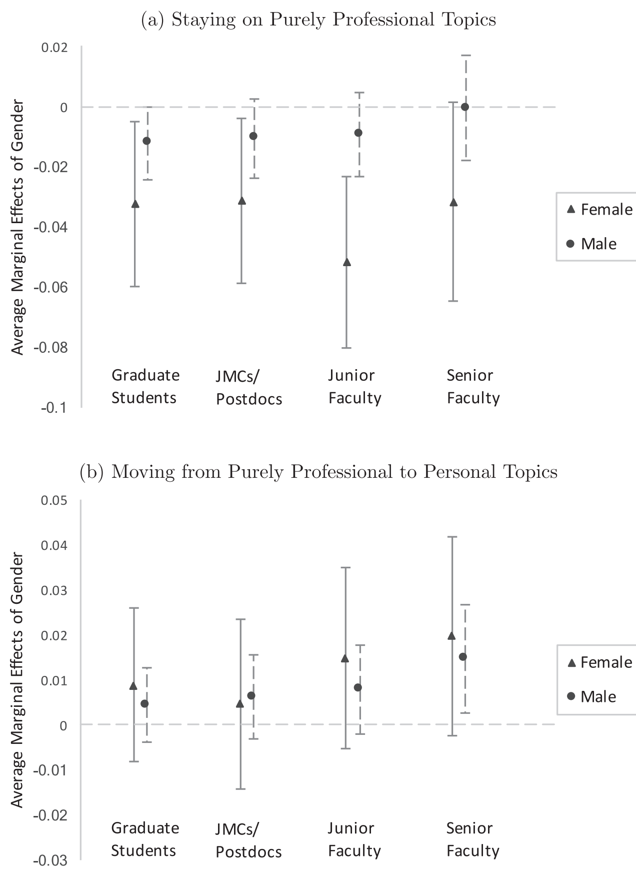

This figure shows the average marginal effects (with 95% confidence intervals) of the mention of a female or a male on the probability of (a) staying on the purely professional topics and (b) moving from professional to personal topics, relative to the baseline where the prior post is not gendered, at four different job ranks. Standard errors are robust and clustered at the thread level.

Although there is less precision in this relatively small sample at each job rank, the results point to alternative hypotheses on posting behavior toward women versus men. First, as discussed in the model of rumors, part of the utility of posting comes from contributing to the public knowledge about the job market (see section II and appendix A). When discussing a specific economist or groups of individuals at a specific job rank, posters may want to get as much professional information as possible. A possible explanation for the contrast between the effects of Female versus Male posts on persistence in professional discussion is that male posters care more about professional discussions of other men than of women. As a result, they are more likely to stay on track in a male-oriented professional discussion. In appendix A, I discuss how the trade-off between contributing to public knowledge and identity boosting leads posters to act differently toward men versus women. However, empirically it is difficult to identify the extent to which the gender gap in movements between topics can be attributed to each incentive separately.

Second, when a poster refers to a specific economist rather than talk about female or male colleagues in general, other posters can form their own interpretations of the subject's professional ability by evaluating his or her work. In this case, the cost of saying something outrageous or deviating to irrelevant personal characteristics is higher to the poster if he takes into account how other posters may react, a component of poster's utility that I discuss in the model of rumors. However, this hypothesis cannot explain why the effect of a Female post at each job rank on the persistence in professional topics remains significantly negative.

In summary, posters' incentive to protect their own professional identity provides a relatively robust explanation for the divergence in the effects of gender on the transition rates between topics. There are alternative explanations based on posters' preferences over different types of threads or over discussions at different job ranks. In future work, it would be particularly meaningful to quantify the trade-off between the incentive to contribute to the public knowledge about the profession and the incentive to boost one's own identity.

VI. Conclusion

This paper uses anonymous discussions on the Economics Job Market Rumors Forum to study people's true attitudes toward women in the profession, which they are unlikely to openly express in other public settings. Posts that discuss women focus significantly less on their professional characteristics and more on physical appearance and personal information than posts that discuss men. Moreover, in the dynamics of conversation, there is a significant lack of persistence in purely professional topics when the prior post mentions a female.

The model of rumors provides an identity-based interpretation of these findings: posters reinforce the perception of women as outsiders in the economics profession through diminishing their professional image, and by doing so, male posters can improve their in-group identity in the profession relative to women. Discussions about women's personal characteristics also cast doubt on the public's understanding about women's true professional ability, which slows the process to overcome hostility against the underrepresented group and improve integration. There are also alternative explanations to these findings. In future analysis, it would be interesting to quantify the role of identity threat in driving discrimination and look further into the trade-off between the incentive to contribute to the public knowledge and the incentive to boost one's own identity.

The stereotypical gender attitudes revealed on the EJMR forum are most likely not exclusive to the economics profession, but reflect the challenges women are facing in many traditionally male-dominated fields. Understanding people's true gender attitudes is crucial to improving policies aimed at increasing diversity in a profession. There is indeed hope to reduce gender bias by promoting interaction between groups (Dahl et al., 2018), increasing exposure to female leaders (Beaman et al., 2009), or more broadly speaking any mechanism that increases information about the true distribution of the ability of women (Goldin, 2015).

REFERENCES

American Economic Association, "AEA Professional Climate Survey: Main Findings," American Economic Association report (2019).

Aguirregabiria, Victor, and Pedro Mira, "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics* 156 (2010), 38–67.

Akerlof, George A., and Rachel E. Kranton, "Economics and Identity," *Quarterly Journal of Economics* 115 (2000), 715–753.

Basford, Tessa E., Lynn R. Offermann, and Tara S. Behrend, "Do You See What I See? Perceptions of Gender Microaggressions in the Workplace," *Psychology of Women Quarterly* 38 (2014), 340–349.

Bayer, Amanda, and Cecilia Elena Rouse, "Diversity in the Economics Profession: A New Attack on an Old Problem," *Journal of Economic Perspectives* 30 (2016), 221–242.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy* 115 (2007), 588–638.

Beaman, Lori, Raghabendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova, "Powerful Women: Does Exposure Reduce Bias?" *Quarterly Journal of Economics* 124 (2009), 1497–1540.

Blau, Francine D., Peter Brummund, and Albert Yung-Hsu Liu, "Trends in Occupational Segregation by Gender 1970–2009: Adjusting for the Impact of Changes in the Occupational Coding System," *Demography* 50 (2013), 493–494.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, "Stereotypes," *Quarterly Journal of Economics* 131 (2016), 1753–1794.

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri, "Are Referees and Editors in Economics Gender Neutral?" *Quarterly Journal of Economics* 135:1 (2019), 269–327.

Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn, and Wendy M. Williams, "Women in Academic Science: A Changing Landscape," *Psychological Science in the Public Interest* 15 (2014), 75–141.

Cotter, David, Joan M. Hermsen, and Reeve Vanneman, "The End of the Gender Revolution? Gender Role Attitudes from 1977 to 2008," *American Journal of Sociology* 117 (2011), 259–289.

Dahl, Gordon, Andreas Kotsadam, and Dan-Olof Rooth, "Does Integration Change Gender Attitudes? The Effect of Randomly Assigning Women to Traditionally Male Teams," NBER working paper 24351 (2018).

Damore, James, "Google's Ideological Echo Chamber: How Bias Clouds Our Thinking about Diversity and Inclusion" (July 2017).

Eckstein, Zvi, and Kenneth I. Wolpin, "Dynamic Labour Force Participation of Married Women and Endogenous Work Experience," *Review of Economic Studies* 56 (1989), 375–390.

Farajtabar, M., J. Yang, X. Ye, H. Xu, R. Trivedi, E. Khalil, S. Li, L. Song, and H. Zha, "Fake News Mitigation via Point Process Based Intervention," arXiv 1703.07823 (2017).

Farajtabar, M., S. Yousefi, L. Q. Tran, L. Song, and H. Zha, "A Continuous-Time Mutually-Exciting Point Process Framework for Prioritizing Events in Social Media," arXiv 1511.04145 (2015).

Fortin, Nicole M., "Gender Role Attitudes and the Labour-Market Outcomes of Women across OECD Countries," *Oxford Review of Economic Policy* 21 (2005), 416–438.
——— "Gender Role Attitudes and Women's Labor Market Participation: Opting-Out, AIDS, and the Persistent Appeal of Housewifery," *Annals of Economics and Statistics*, no. 117/118 (2015).
Gennaioli, Nicola, and Guido Tabellini, "Identity, Beliefs, and Political Conflict," SSRN (2019), https://ssrn.com/abstract=3300726.
Goldin, Claudia, "A Pollution Theory of Discrimination: Male and Female Differences in Occupations and Earnings" (pp. 313–348), in Leah Boustan, Carola Frydman, and Robert Margo, eds., *Human Capital in History: The American Record* (Chicago: University of Chicago Press, 2015).
Hengel, Erin, "Publishing While Female," working paper (2019).
Keane, Michael, and Kenneth I. Wolpin, "The Career Decisions of Young Men," *Journal of Political Economy* 105 (1997), 473–522.
Lundberg, Shelly, "Report: Committee on the Status of Women in the Economics Profession (CSWEP)," *AEA Papers and Proceedings*, 108 (2018), 704–721.
Rizoiu, M.-A., Y. Lee, S. Mishra, and L. Xie, "A Tutorial on Hawkes Processes for Events in Social Media," arXiv 1708.06401 (2017).
Sarsons, Heather, Klarita Gërxhani, Ernesto Reuben, and Arthur Schram, "Gender Differences in Recognition for Group Work," *Journal of Political Economy*, forthcoming.
Tajfel, Henri, and John C. Turner, "Psychology of Intergroup Relations," in S. Worchel and W. G. Austin, eds., *The Social Identity Theory of Intergroup Behavior* (Chicago: Nelson-Hall, 1986).
Wolfers, Justin, "Evidence of a Toxic Environment for Women in Economics," *New York Times*, August 18, 2017.
Wu, Alice H., "Gendered Language on the Economics Job Market Rumors Forum," *AEA Papers and Proceedings* 108 (2018), 175–179.

Online Appendix to "Gender Bias in Rumors Among Professionals: An Identity-based Interpretation"

Alice Wu

May 2019

Appendix A: A Formal Presentation of the Model of Rumors

This appendix provides a formal representation of the model of rumors in Section 2. There are three propositions that correspond to the main predictions summarized in Section 2.

1 Setup

Given a subject i, assume her group $g_i \in \{F, M\}$, and her current job rank $r_i \in \{1, 2...R\}$ where R is the highest rank on the job ladder shared by both groups, are observable to everyone.

Let α_i denotes i's true professional characteristics, and β_i denote her personal characteristics (e.g., physical apperance/family circumstances). Conditional on (g_i, r_i), assume only the professional characteristics α_i matter for the underlying stochastic matching to the job ladder.

Other people get private signals about i's characteristics, through subjectively evaluating the quality of her research works, for example. Let (a_{ij}, b_{ij}) denote poster j's private signals about the subject i's professional and personal characteristics.

I assume a poster's utility comes from two main sources: (1) contribution to public knowledge about the job market, and (2) perception of own identity relative to the subject of discussion.

1.1 Public knowledge about the Job Market

When a poster j enters a thread about i, he observes signals $(a_{i,-j},\, b_{i,-j})$ revealed by previous posters $(-j)$, from which he deduces that the collective estimates of i's professional ability to be μ_i^0

with noise $\sigma_i^0 = \sigma_i^0(a_{i,-j}, b_{i,-j})$. I assume the noise σ_i^0 depends on both the dispersion of professional $a_{i,-j}$ and that of personal $b_{i,-j}$. That is, discussions of personal characteristics do not influence the mean estimate of i's professional ability, but they add noise to the understanding of i's professional ability and thus decrease the signaling value of the mean estimate μ_i^0.

I model the utility from public knowledge as a function of the distance between μ_i^1, the posterior mean signal about i after j's action and j's private information a_{ij}, and the posterior noise σ_i^1:

$$f(|\mu_i^1 - a_{ij}|, \sigma_i^1)$$

If all signals revealed by posters (n previous posts plus j's) are equally weighted, $\mu_i^1 = \frac{n}{n+1}\mu_i^0 + \frac{1}{n+1}a_{ij}$. I make two assumptions about the function f:

- $f_1 < 0$: each poster would like to bring the public knowledge closer to his or her private information
- $f_2 < 0$: each poster also appreciates accurate knowledge

1.2 Identity and Self-image

Following Akerlof and Kranton (2000), I assume poster j cares about his/her self-perceived professional image relative to the subject i, represented by

$$I_j(\alpha_j, \mu_i^1; \vec{\omega}_j) = \begin{cases} \omega_{j,in} \times (\mu_i^1 - \alpha_j) & g_i = g_j \\ -\omega_{j,out} \times (\mu_i^1 - \alpha_j) & g_i \neq g_j \end{cases}$$

where $\vec{\omega}_j = (\omega_{j,in}, \omega_{j,out})$ with $\omega_{j,in} \geq 0$, $\omega_{j,out} \geq 0$ drawn from the same distribution $G(\cdot)$ are the weights poster j put on his image relative to members of his own group ("in-group") versus those of the opposite group ("out-group"). μ_i^1 is the posterior mean estimate of i's professional characteristics, and α_j is poster j's true professional characteristics.

This representation takes into account a person's preference for achieving a positive image of their own group in contrast with the opposite group. That is, a positive description of someone similar to the poster helps affirm a positive identity of his own, whereas it entails a negative externality in terms of identity threat to the poster when $g_i \neq g_j$.

Relative job ranks $r_i - r_j$ can also enter the utility function if we think posters care more about subjects at similar levels than those much higher up on the job ladder.

Assuming that a poster's utility of discussing subject i is additively separable in these two sources and that he takes into account how future posters perturb his utility, the poster's utility function can be written as:

$$\begin{aligned} U_j = & \underbrace{f(|\mu_i^1 - a_{ij}|, \sigma_i^1)}_{\text{public knowledge}} + \underbrace{I_j(\alpha_j, \mu_i^1; \vec{\omega}_j)}_{\text{identity}} \\ + & \underbrace{\beta \times E_j[\triangle f + \triangle I_j]}_{\text{discounted reaction from future posters}} \end{aligned} \tag{1}$$

where the discount factor $\beta \in [0, 1]$, and the perturbations when future posters react to j and move the collective estimate to μ_i' with noise σ_i' are represented by $\triangle f = f(|\mu_i' - a_{ij}|, \sigma_i') - f(|\mu_i^1 - a_{ij}|, \sigma_i^1)$ and $\triangle I_j = I_j(\alpha_j, \mu_i'; \vec{\omega}_j) - I_j(\alpha_j, \mu_i^1; \vec{\omega}_j)$. Intuitively, when a poster j makes an outrageous remark, other posters may react by contradicting his view, and thus takes away part of the utility j has obtained from participating in the discussion.

For simplicity consider two discrete actions poster j can take:

- Action 1 (Professional): reveal his private signal about subject i's professional characteristics a_{ij} but not personal characteristics b_{ij}. Then $\mu_i^1 = \frac{n}{n+1}\mu_i^0 + \frac{1}{n+1}a_{ij}$ and noise σ_i adjusts accordingly.

- Action 2 (Personal): reveal b_{ij} but not a_{ij}. Then the posterior mean estimate for professional ability is the same as the prior μ_i^0 but $\sigma_i \uparrow$ as personal discussions add noise to the professional portrayal of i.

Let $U_j^{(1)}, U_j^{(2)}$ denote poster j's utility of choosing actions 1 and 2, respectively. For now assume $\beta = 0$.

$$U_j^{(2)} - U_j^{(1)} = \begin{cases} f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'}) + \omega_{j,in} \times (\mu_i^0 - \mu_i^1) & g_i = g_j \\ f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'}) - \omega_{j,out} \times (\mu_i^0 - \mu_i^1) & g_i \neq g_j \end{cases} \tag{2}$$

The poster will choose action 2 (Personal) over action 1 (Professional) if $U_j^2 \geq U_j^1$. The choice of action does not depend on the poster's own professional characteristics but depend on the weights he or she puts on the same group versus the opposite group.

2 Propositions

Proposition 1 (Optimal Actions): Given a subject i and the initial public knowledge (μ_i^0, σ_i^0) in a thread, a poster j with private signal (a_{ij}, b_{ij}) will be indifferent between discussing professional characteristics (Action 1) and discussing personal characteristics (Action 2) at group-specific thresholds (a_{in}^*, a_{out}^*) where $a_{in}^* < \mu_i^0$ and $a_{out}^* > \mu_i^0$. Specifically,

1. If $g_i = g_j$, the poster will choose Action 1 if $a_{ij} > a_{in}^*$ and Action 2 if $a_{ij} < a_{in}^*$.
2. If $g_i \neq g_j$, the poster will choose Action 1 if $a_{ij} < a_{out}^*$ and Action 2 if $a_{ij} > a_{out}^*$.

Proof:

First, note that $|\mu_i^1 - a_{ij}| \leq |\mu_i^0 - a_{ij}|$ as the poster can bring the average opinion closer to his own signal through Action 1 - revealing his private signal about i's professional characteristics a_{ij}. In addition, since the noise σ_i increases when there are discussions about personal characteristics under Action 2, assume the noise is higher under action 2 than 1 $\sigma_i^1 > \sigma_i^{1'}$. The utility from contribution to public knowledge decreases in both the distance and the noise factor $f_1 < 0$ and $f_2 < 0$, so we have $f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'}) < 0$.

(1) When the poster comes from the same group as the subject $g_i = g_j$,

$$U_j^2 \geq U_j^1 \longleftrightarrow \underbrace{-\omega_{j,in}}_{\leq 0} \times (\mu_i^0 - \mu_i^1) \leq \underbrace{f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'})}_{<0} \tag{3}$$

Thus the threshold of professional characteristics denoted by a_{in}^* at which the poster is indifferent satisfies $\mu_i^0 > \mu_i^1 \longleftrightarrow a_{in}^* < \mu_i^0$. Note the left hand side (LHS) is monotonically increasing in a_{ij}. When $a_{ij} > a_{in}^*$, and the poster will choose to discuss professional characteristics rather than personal ones. If $a_{ij} < a_{in}^*$, discussing personal characteristics makes the poster better off.

(2) When the poster comes from the opposite group as the subject $g_i \neq g_j$,

$$U_j^2 \geq U_j^1 \longleftrightarrow \underbrace{\omega_{j,out}}_{\geq 0} \times (\mu_i^0 - \mu_i^1) \leq \underbrace{f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'})}_{<0} \tag{4}$$

Thus the threshold of professional characteristics denoted by a^*_{out} at which the poster is indifferent satisfies $\mu_i^0 < \mu_i^1 \longleftrightarrow a^*_{out} > \mu_i^0$. Note the LHS is monotonically decreasing in a_{ij}. When $a_{ij} > a^*_{out}$, the poster will choose to discuss personal characteristics rather than professional ones. If $a_{ij} < a^*_{out}$, discussing professional characteristics makes the poster better off.

□

Figure A1 provides a graphic illustration of the predictions in Proposition 1. The implication is that the average professional signals a given poster reveals about subjects from the in-group is higher than the average professional signals he reveals about subjects from the out-group. Under the framework with two discrete actions, we will also find posters more likely to discuss personal characteristics of people from the out-group.

Figure A1: Optimal Actions

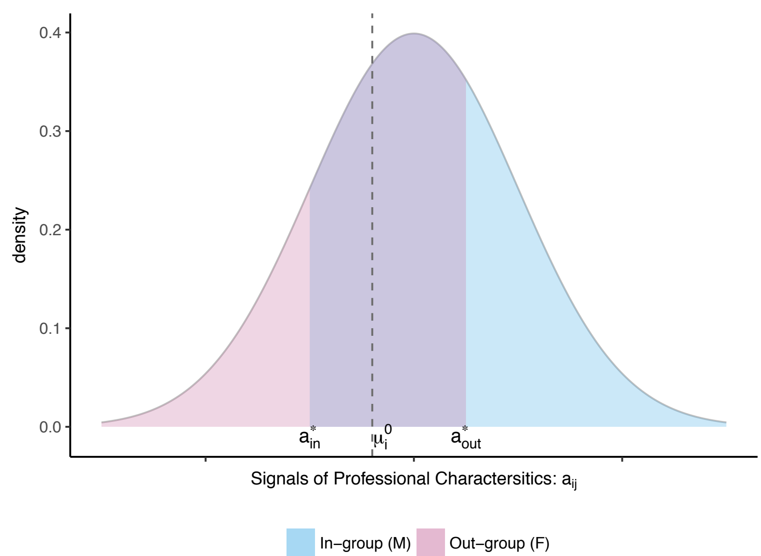

Notes: The shaded region for each group represents the range of professional characteristics of subjects over which the poster is willing to take Action 1 - discussing the subject's professional characteristics but not personal ones. μ_i^0 is the given public knowledge about subject i's professional ability in an existing thread. a^*_{in} is the threshold for the poster's own group above which he will choose Action 1. And a^*_{out} is the threshold for the poster's opposite group below which he will choose Action 1.

Proposition 2 (Heterogeneity in Preferences over Identity): Given a subject i and the initial public knowledge (μ_i^0, σ_i^0) in a thread, the higher the weight a poster puts on his identity relative to a subject from the in-group, $\omega_{j,in}$, the higher the threshold a^*_{in} above which the poster will choose action 1 to reveal his signal about the subject's professional characteristics. On the other hand, the higher the weight a poster puts on his self-image relative one from the out-group, $\omega_{j,out}$, the lower the threshold a^*_{out} above which the poster will choose action 2 to discuss personal rather than professional characteristics.

Proof:

Holding fixed the initial thread with parameters (μ_i^0, σ_i^0), it is immediate from (3) that the threshold a^*_{in} at which the poster is indifferent between action 1 and 2 towards a subject from the in-group is higher as $\omega_{j,in} \uparrow$. Similarly, (4) implies that the threshold a^*_{out} is lower when $\omega_{j,out} \uparrow$.

Figure 2 shows that as both thresholds move closer to the initial μ_i^0, there is a larger gap in the average professional signals a poster is willing to reveal about someone from the in-group versus the out-group.

□

Figure A2: Optimal Actions under Higher Identity Weights

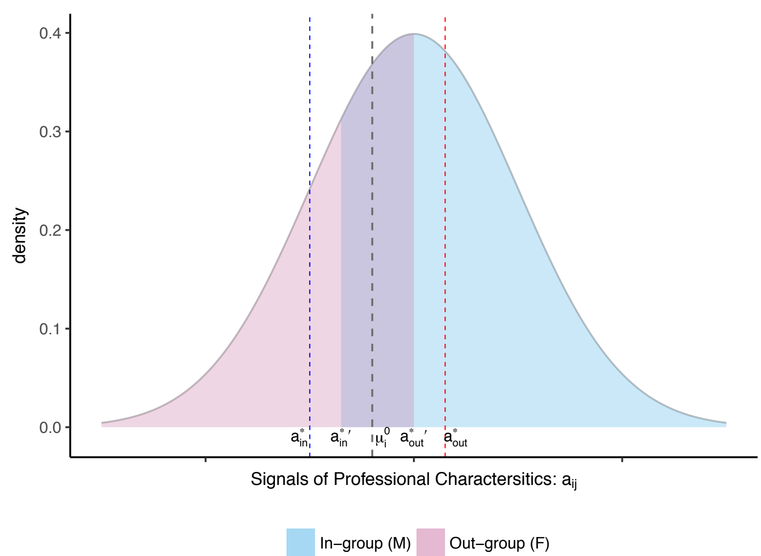

Notes: Relative to the thresholds (a^*_{in}, a^*_{out}), the new thresholds $(a^{*'}_{in}, a^{*'}_{out})$ are for posters who put higher weights on his identity in comparison with the subject of discussion.

Proposition 2 suggests that due to the heterogeneity in the weights posters put on identity, the gap in professional characteristics between group F and M that we observe in the data is a weighted average,

$$E[a_{ij}|i \in M] - E[a_{ij}|i \in F] = \int_{\vec{\omega}} E[a_{ij}|i \in M, \vec{\omega}_j] - E[a_{ij}|i \in F, \vec{\omega}_j] dG(\vec{\omega}_j) \tag{5}$$

Proposition 3 (Interaction between Posters): Given an initial thread with parameters (μ_i^0, σ_i^0), concerns about future reactions ($\beta \in (0,1]$) reduce the fraction of posters who choose to discuss professional rather than personal characteristics of a subject from the opposite group.

Proof:

WLOG, consider the choice of posters with private signals (a_{ij}, b_{ij}) s.t. $a_{ij} > \mu_i^0$ and $b_{ij} > 0$. Such a poster will choose to discuss professional characteristics (Action 1) instead of personal ones (Action 2) if

$$\omega_{j,out} \leq \frac{f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'}) + \beta E_j[\xi]}{\mu_i^0 - \mu_i^1}$$

where ξ represents the perturbation to poster j's utility from other posters reacting differently when j chooses Action 1 vs. Action 2.

If $\beta = 0$, a poster does not care about how future posters react to his action. Then among posters with the same private signals (a_{ij}, b_{ij}), the fraction of posters choosing Action 1 would be:

$$G(\frac{f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'})}{\mu_i^0 - \mu_i^1})$$

where G is the distribution function of weights on identity.

In contrast, if $\beta \in (0,1]$, a poster takes into account a potential backlash in which other posters disagree with him and reduce his expected utility. The fraction of posters with homogeneous private signals choosing Action 1 becomes:

$$G(\frac{f(|\mu_i^0 - a_{ij}|, \sigma_i^1) - f(|\mu_i^1 - a_{ij}|, \sigma_i^{1'}) + \beta E_j[\xi]}{\mu_i^0 - \mu_i^1})$$

which is lower than under $\beta = 0$ if $E_j[\xi] < 0$. Assume a poster cares about how other posters react to his comment on the subject, $\beta \in (0,1]$.

Although I do not have a closed form expression for the reaction component ξ without explicit assumptions about the distributions of private signals and weights, it is reasonable to assume that the magnitude of reaction $|\xi|$ is increasing in the distance between a poster's private signal and the average signals by previous posters, $|\mu_i^0 - a_{ij}|$.

To further illustrate this point on the selection of posters, I use a simple extensive game as in Akerlof and Kranton (2000):

Figure A3: Interactions between Posters

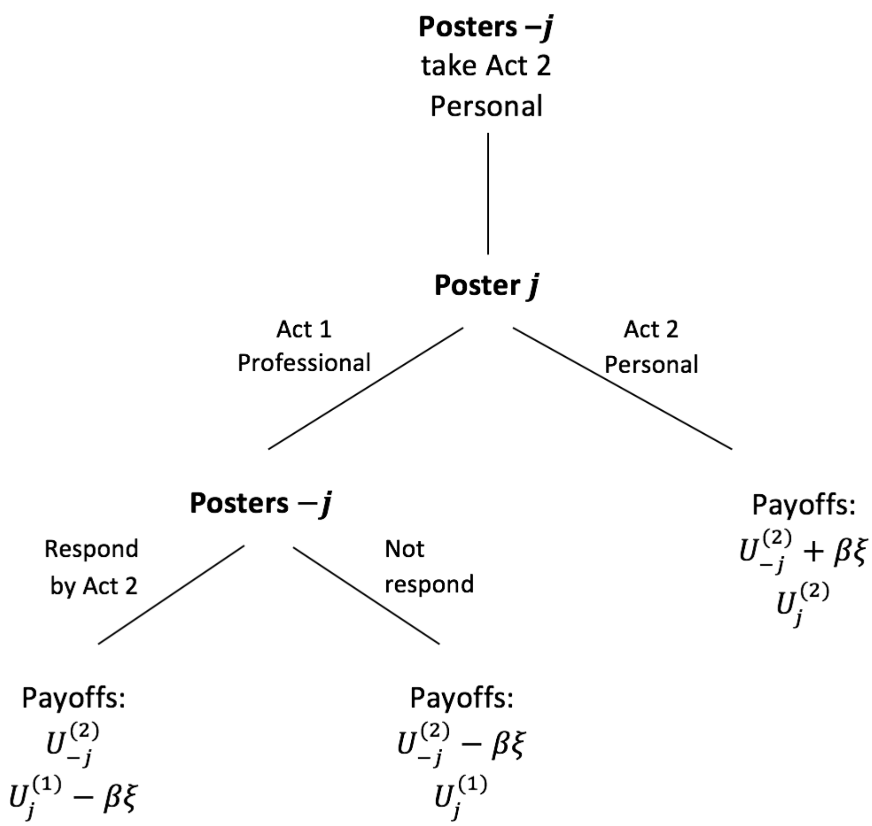

Notes: Assume $\xi > 0$. When a new poster j enters the thread, he observes that other posters $-j$ take Action 2 to emphasize personal characteristics of the subject. j may want to talk about i's professional characteristics (Action 1) instead and gets $U_j^{(1)} > U_j^{(2)}$. However, other posters $-j$ give a credible threat when $\beta > 0$. As a result, if $U_j^{(1)} - \beta\xi < U_j^{(2)}$, poster j will comply with other posters by taking action 2 or avoid posting.

Appendix B: Data

This appendix provides additional information about the sample construction.

1. Web Scraping

I used the Beautiful Soup package on Python to scrape data from the forum.[1] There are two steps in my scraping process. First, I went to each page of the forum to obtain the title of each thread and its URL, and basic information about the thread such as the number of posts and a rough time stamp for its latest update. Figure B1 shows a screenshot of the main site of the EJMR forum.

Second, I entered each thread initiated or updated between October 2013 and October 2017 to access the posts on its first and last page.[2] Figure B2 provides a screenshot of a EJMR thread and posts on its first page. About 88% of threads in these four-year sample contain only one page with ≤ 20 posts.

The reason for scraping only the first and the last page of each thread is that some particular threads contain over a thousand pages (e.g., a thread titled "German Market" has about 1,700 pages), and users often browse through the first page to understand the initial topic of a thread, and then go to the last page to see the latest update of a thread and decide whether to join the thread or not. Therefore, the first and the last page of each thread provides a reasonable summary of the beginning and the latest update/conclusion of the thread.

2. Gendered Posts

I identify Female and Male posts through gender classifiers that arise from the most frequent 10,000 words, and the names of about 9,000 economists. Table B1 lists all classifiers that indicate a Female or a Male post. Table B2 lists the number of female and male classifiers, and female and male economists in two datasets I assembled.

[1] The documentation of the Beautiful Soup package can be found at:
https://www.crummy.com/software/BeautifulSoup/bs4/doc/

[2] The time stamps do not allow me to identify the month in which a thread is created beyond one year. I use "4 years ago" as a indicator for threads last updated around October 2013, four years ago from the last round of scraping in 2017.

- **NBER Authors**: the EJMR forum has been posting abstracts of new NBER working papers every week since 2014. I assembled a dataset of authors of $4,478$ NBER working papers from 2014 to 2017. There are 5,003 authors in total, among whom 1,008 are affiliated with NBER as research associates, 301 are faculty research fellows and the rest are their collaborators. This sample represents active researchers in economics.

- **Job Market Candidates**: I focused on economics departments from the U.S. or Canada that rank among top 50 in the global ranking by econphd.net 2004. I found names of Ph.D. graduates through two sources:

  - ProQuest Disseration Database: I searched for doctoral dissertations in economics from each school on my list. I was able to find dissertations for 31 schools out of 37 on my list.

  - Placement Records Online: I went to the department websites if I couldn't find any record on ProQuest or if the number of Ph.D. graduates fluctuated too much by year. I found placement records with names of job market candidates on the websites of 11 economics departments.

  From these two sources I found 4,724 Econ Ph.D. who graduated from 36 top programs between 2011 and 2018.[3]

Table B3 provides a list of departments, and number of female and male candidates I found. Please note that the sample of job market candidates does not mean to be a complete list of Ph.D. graduates from all 36 economics departments. It serves the purpose to identify additional gendered posts that refer to specific economists.

I assign gender to the 9,000 economists in these two samples through:

- A dataset of economists with gender assignment, assembled by Card, DellaVigna, Funk and Iriberri (2018): this dataset includes about 48,000 authors with a published article between 1990 and mid-2017 in a set of 53 economics journals. They assigned gender to authors

[3] There is only one department for which I cannot find records of job market candidates by this method. Due to data protection policy of the school's administrative data, I decided not to proceed with the data application process.

through running automatic packages in R on first names, matching with additional lists of female economists and manually searching names online.

- Packages that assign gender based on first names: "gender" and "genderizeR" on R[4]
- Manual assignment for about 900 economists

By gender classifiers alone, I identified 102,956 Female posts and 327,670 Male posts. About 87% of posts that include economists' names have already been picked up by gender classifiers, as these posts are likely to include pronouns like "he" or "she". In total, I found 104,476 Female posts and 334,721 Male posts, comprising over 20% of all posts during the sample period. I preserve all threads that include at least one Female or Male post and call it the gender sample, which contains 1.7 million posts in total.

3. Topics

Given the 2.2 million posts in my four-year dataset, I identified the most frequent 10,000 words and then manually classified them into 15 categories. Table B4 shows how I group these categories into two main topics: Academic/Professional, and Personal/Physical.

4. Job Ranks

About 15% of posts in the gender sample (1.7 million posts in total including about 100,000 Female and 330,000 Male posts) can be assigned one of the four job ranks: Graduate Students, JMCs/Post-docs, Junior Faculty and Senior Faculty. First, I use information about economists in my sample of NBER authors and sample of job market candidates to classify their job rank. For job market candidates with years of graduation from Ph.D., I assigned posts in threads created more than one year before graduation to the level of graduate students, and those ±1 year relative to graduation to the level of JMC/post-docs, and those after to the level of junior faculty. For NBER authors who do not occur in the sample of job market candidates, I use their job ranks as of 2016.

[4]See documentations at https://cran.r-project.org/web/packages/gender/gender.pdf and https://github.com/kalimu/genderizeR

Finally, I use keywords indicating a job rank to identify additional posts at each rank. Please see Table B5 for the lists. When a post contains keywords from multiple ranks, I assign the highest rank possible.

Figure B1: Screenshot of the First Page of the EJMR Forum

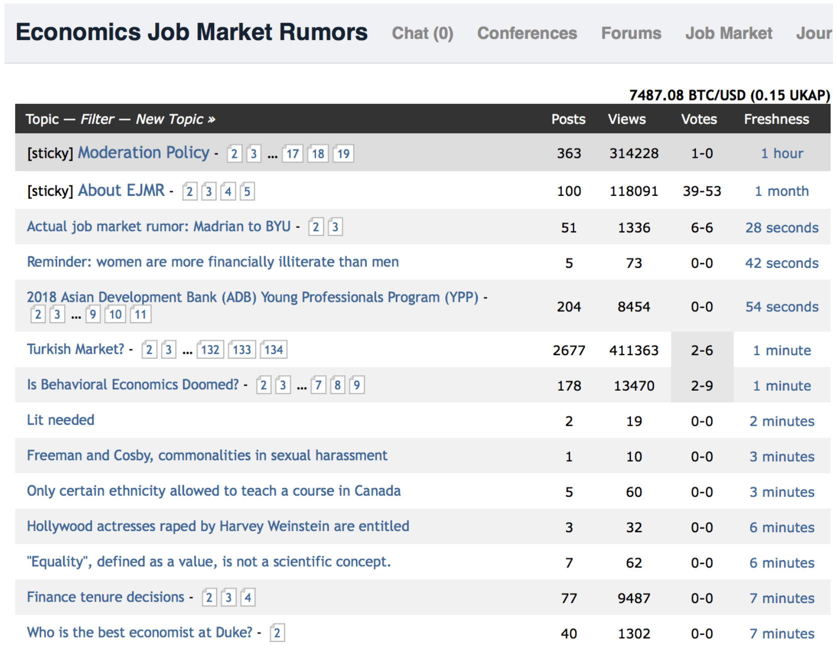

Economics Job Market Rumors Chat (0) Conferences Forums Job Market Jour

7487.08 BTC/USD (0.15 UKAP)

Topic — *Filter* — *New Topic »*	Posts	Views	Votes	Freshness
[sticky] Moderation Policy - 2 3 ... 17 18 19	363	314228	1-0	1 hour
[sticky] About EJMR - 2 3 4 5	100	118091	39-53	1 month
Actual job market rumor: Madrian to BYU - 2 3	51	1336	6-6	28 seconds
Reminder: women are more financially illiterate than men	5	73	0-0	42 seconds
2018 Asian Development Bank (ADB) Young Professionals Program (YPP) - 2 3 ... 9 10 11	204	8454	0-0	54 seconds
Turkish Market? - 2 3 ... 132 133 134	2677	411363	2-6	1 minute
Is Behavioral Economics Doomed? - 2 3 ... 7 8 9	178	13470	2-9	1 minute
Lit needed	2	19	0-0	2 minutes
Freeman and Cosby, commonalities in sexual harassment	1	10	0-0	3 minutes
Only certain ethnicity allowed to teach a course in Canada	5	60	0-0	3 minutes
Hollywood actresses raped by Harvey Weinstein are entitled	3	32	0-0	6 minutes
"Equality", defined as a value, is not a scientific concept.	7	62	0-0	6 minutes
Finance tenure decisions - 2 3 4	77	9487	0-0	7 minutes
Who is the best economist at Duke? - 2	40	1302	0-0	7 minutes

This figure shows what the main page of the EJMR forum looks like. Users browse through the titles of threads before accessing the posts in each thread. The main pages also record total number of posts, number of views, votes by users, and a rough time stamp of the latest post in each thread.

Figure B2: Screenshot of a EJMR Thread

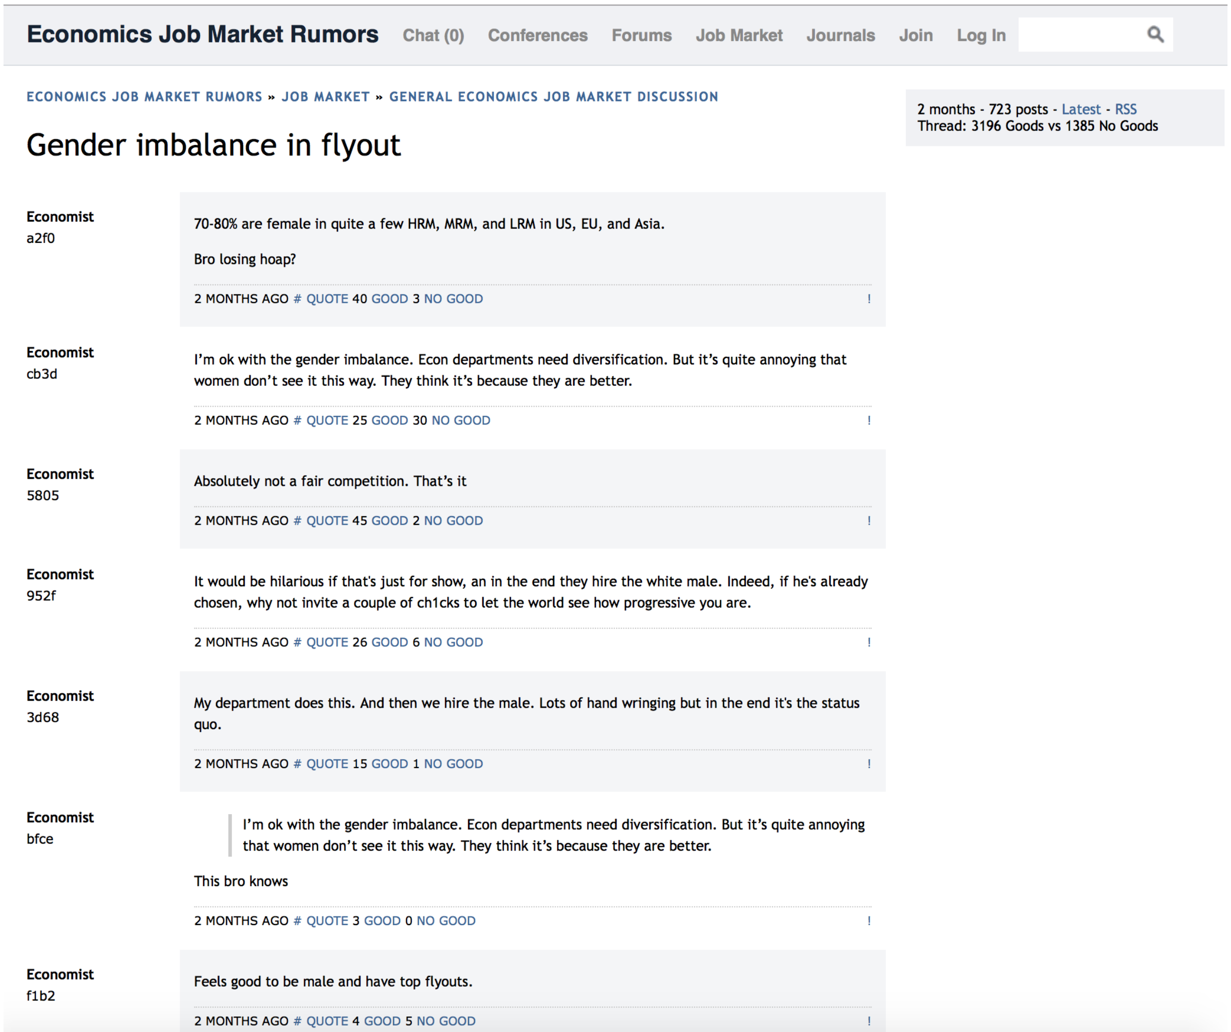

This figure shows what a thread looks like on EJMR. Once users click on a thread title from the main site (see Figure B1), they enter the thread environment and observe the posts under the thread, arranged in chronological order.

Table B1: Comprehensive List of Gender Classifiers from the Most Frequent 10,000 Words

	Indicating *Female* Posts	Indicating *Male* Posts
Pronouns	her, herself, she, she'd, she'll, she's	he, he'd, he'll, he's, him, himself, his
Names	amy athey duflo elizabeth emily hilary hillary jane jennifer jessica maria mary nancy reinhart sarah susan yellen	acemoglu adam akerlof alan albert alexander allen andrew angrist arthur autor baker barro becker ben benjamin bernanke berry blanchard bob borjas brian campbell carl carlos charlie chetty chris christopher cochrane colin cowen daniel daron david deaton delong duffie edward eric eugene fama frank frey friedman friedman's gary george gintis glaeser gordon greene greg gregory hansen harry hayashi hayek heckman henderson henry imbens jack james jason jeff jeffrey jeremy jimmy joe john jon jonathan jose joseph justin kehoe ken kenneth kevin krueger kruggles krugman krugman's larry lars levine levitt lucas mankiw mark martin matt matthew michael mike miller milton murphy myerson neumann nicholas nick noah parag pat pathak patrick paul perez peter phil philip phillips pierre piketty piketty's powell prescott rabin raj ravikumar ricardo richard robert roberts robinson roger rogoff ron roth rubin rubinstein russ rust ryan saez sam sargent shapiro shiller shleifer simon sims stephen steve steven stiglitz summers terry thaler thomas tim tirole tom tony victor walker wallace walters werning williamson wolfers woodford wooldridge
General Identity	female, females, ladies, lady, woman, women, women's	male, males, man, man's, men, men's
Miscellaneous	bietches bitch bitches broette broettes chick chicks daughter daughters gf girl girlfriend girlfriends girls mom moms mother mothers sister sisters wife wife's wives	bf boy boyfriend boys bro bros brother brothers dad daddy dude dudes father fathers gentleman grandfather guy guy's guys husband husbands papa sir son sons uncle

Note: All of the gender classifiers emerge among the most frequent 10,000 words from the raw 2.2 million posts in the four-year sample. "Names" are either first names and economists' last names from which I can tell the gender. Gender neutral first names or last names that I cannot relate to specific economists are not included.

Table B2: Identifying Gendered Posts

	Female	Male
I. Gender Classifiers		
Pronouns/Group Identities	36	39
Names	17	165
II. NBER Authors		
All	1,215	3,788
Discussed on EJMR	366	1,654
III. Job Market Candidates		
All	1,482	3,242
Discussed on EJMR	221	650

Notes: Table B1 gives the complete list of gender classifiers. NBER Authors consist of authors of NBER working papers between 2014 and 2017. Job Market Candidates consist of recent Economics Ph.D. graduates between 2011 and 2018.

Table B3: Schools in the Sample of Job Market Candidates

School	Total	Females	Males
Berkeley	206	70	136
Boston College	68	22	46
Boston University	129	51	78
British Columbia (UBC)	60	20	40
Brown	66	20	46
CalTech	16	2	14
Carnegie Mellon (CMU)	61	17	44
Columbia	192	72	120
Cornell	128	43	85
Duke	185	62	123
Harvard	310	94	216
Indiana	87	32	55
Maryland - College Park	145	51	94
Michigan - Ann Arbor	215	64	151
Michigan State	95	30	65
Minnesota	137	49	88
MIT	170	58	112
North Carolina - Chapel Hill	97	36	61
Northwestern	128	24	104
NYU	114	21	93
Ohio State (OSU)	119	45	74
Princeton	147	33	114
Penn State (PSU)	118	32	86
Rochester	67	22	45
Stanford	161	44	117
Texas A&M	129	34	95
UChicago	218	63	155
UCLA	138	40	98
UCSD	110	38	72
UC Davis	142	44	98
UIUC	83	26	57
UPenn	179	53	126
USC	85	29	56
UToronto	82	32	50
Wisconsin - Madison	168	60	108
Yale	169	49	120
Total	4,724	1,482	3,242

Note: This table shows the number of Economics Ph.D. graduates in my sample by school. I assembled the data from two sources: (1) ProQuest dissertation database, and (2) placement records of job market candidates on the websites of economics departments.

Table B4: Categories of Words

Category	No. Words	Examples
All Gender Classifiers		
Female	53	“she”, “female”
Male	204	“he”, “male”
Academic/Professional		
Economics	140	“economics”, “macro”, “empirical”,“QJE”, “Keynesian”
Academic-General	1,295	“research”, “papers”,“tenure”, “teaching”, “professor”
Professional	180	“career”, “interview”, “payrolls”, “placement”, “recruit”
Personal/Physical		
Personal Information	113	“family”,“married”, “kids”, “relationship”,“lifestyle”
Physical Attributes	134	“beautiful”, “handsome”, “attractive”, “body”,“fat”
Gender related	67	“gender”, “feminine”, “masculine”, “sexist”, “sexual”
Swear Words		
Swear	67	“shit”, “wtf”, “asshole”
Intellectual		
Intellectual-Positive	106	“intelligent”, “creative”,“competent”
Intellectual-Neutral	32	“brain”, “iq”, “ability”
Intellectual-Negative	124	“dumb”,“ignorant”,“incompetent”
Miscellaneous		
Emotion/Feelings	121	“happy”,“depressing”
Others	7,431	“years”, “places”, “everything”
Total	10,000	

Notes: “Gender related” category under *Personal/Physical* are not used as gender classifiers.

Table B5: Identifying Job Rank by Keywords in Each Post

Job Rank	Keywords
Graduate Students	“research assistant”, “ra” (RA), “graduate student”, “grad student”, “phd”, “ta” (TA), “cohort”, “classmate”, “colleague”, “coauthor”, “co author”, “office mate”, “officemate”
Job Market Candidates / Post-docs	“candidate”, “job market”, “jmc” (job market candidate), “jmp”(job market paper), “placement”, “flyout”, “post-doc”, “post doc”, “post-doc”
Junior Faculty	“junior faculty”, “assistant professor”, “assistant prof”, “ap” (Assistant Professor), “associate professor”, “associate prof”, “tenure”, “untenured”, “tenured”, “midterm review”
Senior Faculty	“full professor”, “full prof”, “chaired”, “endowed prof”, “endowed chair”, “senior faculty”, “department chair”, “editor”, “nobel”, “bates clark”, “clark prize”, “clark medal”, “fischer black prize”

Notes: This table lists the keywords I use to determine the job rank among posts that do not mention an economist with job rank assigned yet. If a post includes keywords from multiple ranks, I assign the highest job rank possible.