AEA Papers and Proceedings 2018, 108: 175–179
https://doi.org/10.1257/pandp.20181101

GENDER ISSUES IN ECONOMICS

Gendered Language on the Economics Job Market Rumors Forum†

By Alice H. Wu*

Women are underrepresented in math-intensive fields (Ceci et al. 2014; Kahn and Ginther 2017), and analysts have noted that the representation gap is as large or larger in economics than in STEM (science, technology, engineering, and math) fields on average (e.g., Bayer and Rouse 2016). Among various mechanisms that have been proposed to explain this gap,[1] one that seems particularly relevant but that has not yet been evaluated systematically, is the role of an unwelcoming culture that reinforces stereotypical beliefs of men as an in-group in the field and women as an out-group (e.g., Tajfel and Turner 1986; Tonso 1996).

This paper attempts to assess the existence of an unwelcoming or stereotypical culture using evidence on how women and men are portrayed in anonymous discussions on the Economics Job Market Rumors forum (EJMR). As its name suggests, EJMR was established to share information about job interviews and outcomes in each year's hiring cycle, though it is active year-round. EJMR users post anonymously about economics-related or miscellaneous issues. Anonymity presumably eliminates social pressures that constrain participants' speech in other public settings, leading to a record of postings that reveal what participants believe but would not otherwise openly express.

I use a Lasso logistic model to measure gendered language in EJMR postings, identifying the words that are most strongly associated with discussions about one gender or the other. I find that the words most predictive of a post about a woman (*female* words) are generally about physical appearance or personal information, whereas those most predictive of a post about a man (*male* words) tend to focus on academic or professional characteristics. Despite some intervention by EJMR moderators, the top *female* words include several explicitly sexual terms. Gendered language is also shown to be widespread: about one in five posts about women (*Female* posts) contains at least one of the top 50 *female* words selected by Lasso, many of which are arguably inappropriate for a professional forum. Finally, I evaluate the robustness of the word-selection process through a subsampling exercise, which provides more confidence in the conclusion of differential portrayal of women and men on the forum.

I. Data

I scraped 2,217,046 posts on the first and last page of 223,475 threads on EJMR initiated or updated between October 2013 and October 2017. In the absence of a pre-existing dictionary, I identified the most frequent 10,000 words from the raw text and recorded the word counts for each word in each post. To determine the gender of the subject of each post, I extracted a list of 57 female classifiers (e.g., "she"/"woman") and a list of 236 male classifiers (e.g., "he"/"man")

* University of California, Berkeley (email: haowen.wu@berkeley.edu). I am extremely grateful to David Card and Janet Currie for their guidance and support. I also thank David Romer, Patrick Kline, Amanda Pallais, Jesse Rothstein, Michael Reich, Will Dobbie, Justin Wolfers, Alessandra Fenizia, Sydnee Caldwell, Ulrike Malmendier, Wei Jiang, Dick Startz, Patrick Button, Jonathan Tan, seminar participants at UC Berkeley, Harvard Business School, and Bowdoin College for helpful comments.

† Go to https://doi.org/10.1257/pandp.20181101 to visit the article page for additional materials and author disclosure statement.

[1] For example, recent studies examine course-taking patterns and comparative advantage (Card and Payne 2017), the impacts of role models (Carrell, Page, and West 2010), and stereotype beliefs (Reuben, Sapienza, and Zingales 2014; Bordalo et al. 2016).

from the top 10,000 words. The gap between the numbers of female and male classifiers is driven by the different numbers of female and male names among the top 10,000 words.

I consider a post to be *Female* if it contains any female classifier and *Male* if it contains any male classifier. Using the comprehensive list of gender classifiers, I identify 444,810 gendered (*Female* or *Male*) posts, comprising over 20 percent of all posts over the past four years. These gendered posts are from 138,477 threads, representing about 62 percent of all threads in the past four years.

II. Lasso Logistic Model

I fit a Lasso logistic model to predict whether a gendered post is *Female* or *Male* using the types of words in the post. My hypothesis is that an unwelcoming or stereotypical culture will lead EJMR participants to use terms to describe men that emphasize their fit and position within the field and terms to describe women that de-emphasize their professional accomplishments. Specifically, letting $\mathbf{w}_i$ denote a vector of counts for each of the most common words (excluding all gender classifiers) that are present in gendered post i, I estimate a Lasso-regularized logistic model for the probability that the post is *Female*, as follows:

$$\hat{\theta}_\lambda = \arg\min_\theta - \log(\Pi_{i=1}^N P(Female_i \mid \mathbf{w}_i)) + \lambda ||\theta||_1,$$

where $\|\theta\|_1 = \sum_{j \geq 1} |\theta^j|$. The Lasso regularization helps identify words with the strongest predictive power while avoiding over-fitting. I estimate the model using only gendered posts that refer uniquely to one gender or the other, excluding posts that contain classifiers for both genders (which account for about 10 percent of gendered posts).

A. *Model Training Process*

There are 401,734 posts that include only female or only male classifiers from the comprehensive list. I use a 75 percent random sample to train the model and select an optimal tuning parameter λ^* through 5-fold cross validation. I then select the p-score threshold that minimizes the mean squared error for predicting gender on the remaining 25 percent as the test set.[2] This leads to the selection of a threshold of $p^* = 0.40$ for assigning a post to be *Female*. I use the same threshold to assign genders for the posts that include both female and male classifiers: 31.8 percent of the posts that contain classifiers for both genders are re-classified to *Female*, and the rest to *Male*.

B. *Gendered Words*

The estimated model identifies about 4,500 words with nonzero predictive power for determining gender. I sorted these words by their marginal effect, i.e., the increase in the probability that the subject of a post is *Female* given an additional occurrence of each word. Table 1 displays the top ten words selected by Lasso.

The table reveals that the words that are most predictive of a *Female* post are typically about: (i) physical appearance; (ii) personal or family information; or (iii) gender issues/sexism. The words "hot" and "attractive" increase the predicted probability that a post is discussing a female by approximately 27.1 percent and 24.5 percent, respectively. While such terms could be viewed as complimentary in other settings, in this setting they arguably reflect the treatment of women as an out-group who are to be judged by nonprofessional standards (e.g., physical appearance). For example, there is a thread titled "Cute, unmarried HRM AP is doing a seminar at my school. Can I ask her out?",[3] which judges a female economist based on her appearance, with no reference to professional-related attributes.

In contrast the words that predict a *Male* post include more academically and professionally oriented terms. For example, "adviser," "supervisor," and "Nobel" are in the 30 most predictive *male* words, and each increases the probability that a post is discussing a male by about 13 percent–15 percent. Nevertheless, the Lasso model also picks up a few offensive (and potentially out-group-defining) terms such as "homo," suggesting an unwelcoming online environment for some subgroups of males.

[2] See online Appendix Figure 1 for a plot of MSE at each p-score cutoff.

[3] This thread was initiated and last updated 2 years ago. It contains 20 posts and has 1, 238 views.

Table 1—Top 10 Words Most Predictive of Female/Male

Most *female*		Most *male*	
Word	ME	Word	ME
Hotter	0.422	Homo	−0.303
Pregnant	0.323	Testosterone	−0.195
Plow	0.277	Chapters	−0.189
Marry	0.275	Satisfaction	−0.187
Hot	0.271	Fieckers	−0.181
Marrying	0.260	Macroeconomics	−0.180
Pregnancy	0.254	Cuny	−0.180
Attractive	0.245	Thrust	−0.169
Beautiful	0.240	Nk	−0.165
Breast	0.227	Macro	−0.163

Notes: The model was trained on a 75 percent sample of gendered posts that contain only female or only male classifiers from the comprehensive list. ME—the marginal effect of word w is the change in probability that a post is discussing a female, when it contains an additional word w. The words that predict *Female* (*Male*) are sorted in descending (ascending) order of the ME.

The moderation policy on the EJMR forum is based both on an automatic censorship of words and on reports by users. The evidence of stereotyped and offensive language captured here suggests that either the moderators did not remove the threads reported by users, or that the users themselves tolerated such content and did not complain.

To make inferences about the pervasiveness of gendered language, I consider the frequency of the words selected by Lasso.[4] Some of the most *female* words also turn out to be relatively common. For example, the word "hot" shows up in about 3.5 percent of the *Female* posts, and ranks as the third most common term in *Female* posts, whereas the third most common word in *Male* posts is "job." Overall, about 19.4 percent of all *Female* posts include at least one of the top 50 *female* terms, most of which highlight physical attributes or personal information.

III. Robustness Check

One concern about assigning gender to posts based on the comprehensive list of gender classifiers is that it may over-identify gendered words that occur in personal discussions about "girlfriends" or "boyfriends," which are included as classifiers. To address this concern, I conduct a robustness check by replicating the analysis using gendered posts identified only by gender pronouns (e.g., "he" or "she"). Relative to gendered posts identified using the comprehensive list, more of the posts identified using pronouns, referred to as the pronoun sample, pertain to specific individuals. As a result, the model trained on the pronoun sample should pick up more academic or professional terms for both genders.

Following the same procedure, I train a Lasso logistic model[5] on 35,850 *Female* posts and 103,449 *Male* posts in the pronoun sample. As expected, the estimated model based on the pronoun sample identifies a few more academic terms. For example, "AEJ" (ME: 13.6 percent) and "RCT" (ME: 13.3 percent) appear among the top *female* words.[6] The marginal effects of terms such as "advisor," "Nobel," and "promoted" among the top *male* words become stronger. Nevertheless, an overwhelming majority of the *female* words continue to focus on non-academic aspects. For example, six out of the ten most *female* words selected when gender is determined by the comprehensive list of classifiers also appear when it is determined by pronouns only (see Table 2).

Finally, to evaluate the pervasiveness of gendered words identified using the two alternative sets of classifiers, Figure 1 plots the fraction of *Female* (*Male*) posts that contain at least one of the 50 words most strongly associated with *Female* (*Male*) under the two alternatives.

This trend plot for data in the most recent year reveals several interesting patterns of gendered language. First, there is a large gap between the pervasiveness of the top *female* versus top *male* words selected by Lasso, particularly when gendered posts are identified using the comprehensive list of gender classifiers. Across all months, about 17.2 percent to 19.6 percent of all *Female* posts identified by the comprehensive list include at least one of the top 50 *female* words, but for *male* words the equivalent measures

[4] See online Appendix Table 1, 2, and 3 for the top 50 *female* and *male* words selected by Lasso, the number of gendered posts each of the words occurs in, and the most frequent 50 words in gendered posts, respectively.

[5] For an additional check, I train a Lasso-regularized linear probability model on the pronoun sample, and the top 50 *female* or *male* words selected by the linear Lasso are shown in online Appendix Figure 2 and online Appendix Figure 3.

[6] For word selection by Lasso logistic on the pronoun sample, see online Appendix Tables 4, 5, and 6.

Table 2—Top 10 Words Most Predictive of Female/Male (*Pronoun sample*)

Most *female*		Most *male*	
Word	ME	Word	ME
Pregnancy	0.292	Knocking	−0.329
Hotter	0.289	Testosterone	−0.204
Pregnant	0.258	Blog	−0.183
Hp	0.238	Hateukbro	−0.176
Vagina	0.228	Adviser	−0.175
Breast	0.220	Hero	−0.174
Plow	0.219	Cuny	−0.173
Shopping	0.207	Handsome	−0.166
Marry	0.207	Mod	−0.166
Gorgeous	0.201	Homo	−0.160

Note: The model was trained on a 75 percent sample of gendered posts that contain only feminine pronouns or only masculine pronouns.

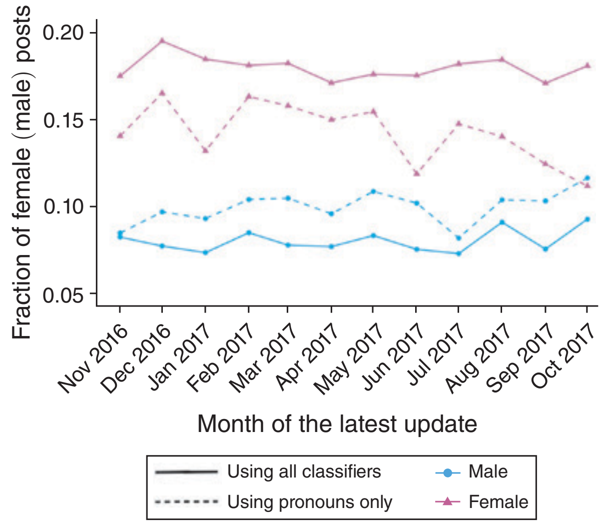

Figure 1. Fraction of *Female* (*Male*) Posts that Include any Top 50 *Female* (*Male*) Words, under Two Alternatives

Notes: The solid lines plot the fraction of *Female* (*Male*) posts identified by the comprehensive list of gender classifiers that include at least one of the top 50 *female* (*male*) words selected by the Lasso-Logistic model. The dashed lines plot the equivalent measures for the word selection based on gendered posts identified by pronouns only. For threads initiated or updated from November 2016 to October 2017, I identified the month of its most recent post from the rough time stamps on the main pages of EJMR.

range from 7.3 percent to 9.3 percent. In the pronoun sample, the gap in pervasiveness shrinks: the top *female* words selected when gender is identified only by pronouns become less common, whereas the top *male* words become more common.

Second, there is larger month-to-month variation in the pervasiveness of the top *female* words selected through the pronoun sample than through the complete sample, especially during the job market season. It is disturbing to see that within the pronoun sample, the fraction of *Female* posts that include at least one of the top 50 *female* terms can be 3 to 4 percentage points higher in particularly active months of the job market (December 2016, February and March 2017) than other months. Such variation suggests that the competitive environment of the job market may lead to more gendered discussions about female and male candidates.

Third, there is evidence of some effect of media discussions about the content of EJMR postings in August 2017. A *New York Times* article by Justin Wolfers,[7] citing results from Wu (2017), raised some concerns about the gendered discussions on EJMR. This treatment appears to have led to a decline in the occurrences of the top *female* words in the pronoun sample in the following two months, which may reflect either a decrease in the usage of gendered words or an increase in censoring by EJMR moderators. If the censoring is playing a more important role, however, then this trend should not be interpreted as a change in the underlying beliefs or attitudes of the posters.

To summarize, despite the differences in the pervasiveness of gendered words selected under the two alternatives, this robustness check confirms that the postings about women tend to highlight physical appearance, personal information, and sexism, whereas those about men are more academically or professionally oriented.

IV. Discussion

This paper illustrates the use of text analytic techniques to measure gendered language between posts pertaining to women and men. The gendered posts may not necessarily talk about specific female or male academics, but they play a large role in shaping the overall

[7] Wolfers, Justin. 2017. "Evidence of a Toxic Environment for Women in Economics." *New York Times*, August 18. https://www.nytimes.com/2017/08/18/upshot/evidence-of-a-toxic-environment-for-women-in-economics.html.

atmosphere on this forum for economists, which may consolidate the perception of men as an in-group versus women as an out-group.

However, an analysis at the word level provides an incomplete picture of the stereotyping behavior on EJMR. Wu (2017) designs a topic analysis and provides an econometric framework for quantifying stereotyping in the dynamics of conversation. Wu (2017) also shows that high-profile female economists tend to receive more attention than their male counterparts, which may suggest that the work by women is more heavily scrutinized.

REFERENCES

Bayer, Amanda, and Cecilia Elena Rouse. 2016. "Diversity in the Economics Profession: A New Attack on an Old Problem." *Journal of Economic Perspectives* 30 (4): 221–42.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2016. "Stereotypes." *Quarterly Journal of Economics* 131 (4): 1753–94.

Card, David, and A. Abigail Payne. 2017. "High School Choices and the Gender Gap in STEM." National Bureau of Economic Research Working Paper 23769.

Carrell, Scott E., Marianne E. Page, and James E. West. 2010. "Sex and Science: How Professor Gender Perpetuates the Gender Gap." *Quarterly Journal of Economics* 125 (3): 1101–44.

Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn, and Wendy M. Williams. 2014. "Women in Academic Science: A Changing Landscape." *Psychological Science in the Public Interest* 15 (3): 75–141.

Kahn, Shulamit, and Donna Ginther. 2017. "Women and STEM." National Bureau of Economic Research Working Paper 23525.

Reuben, Ernesto, Paola Sapienza, and Luigi Zingales. 2014. "How Stereotypes Impair Women's Careers in Science." *Proceedings of the National Academy of Sciences of the United States of America* 111 (12): 4403–08.

Tajfel, Henri, and John C. Turner. 1986. "The Social Identity Theory of Intergroup Behavior." In *Psychology of Intergroup Relations*, edited by William G. Austin and Stephen Worchel, 7–24. Chicago: Nelson-Hall.

Tonso, Karen L. 1996. "The Impact of Cultural Norms on Women." *Journal of Engineering Education* 85 (3): 217–25.

Wu, Alice H. 2017. "Gender Stereotyping in Academia: Evidence from Economics Job Market Rumors Forum." Unpublished.

Online Appendix to "Gendered Language on the Economics Job Market Rumors Forum"

Alice H. Wu

Model I. Lasso-regularized Logistic Model

Letting $\mathbf{w}_i$ denote a vector of counts for each of the most common words (excluding the female or male classifiers) that are present in gendered post i, I assume the posterior probabilities are:

$$\begin{aligned} P(Female_i = 1|\mathbf{w}_i) &= \frac{exp(\theta_0 + \mathbf{w}_i'\theta)}{1 + exp(\theta_0 + \mathbf{w}_i'\theta)} \\ P(Female_i = 0|\mathbf{w}_i) &= \frac{1}{1 + exp(\theta_0 + \mathbf{w}_i'\theta)} \end{aligned} \tag{1}$$

Write the likelihood of each observation as:

$$P(Female_i|\mathbf{w}_i) = P(Female_i = 1|\mathbf{w}_i)^{Female_i} \times P(Female_i = 0|\mathbf{w}_i)^{(1-Female_i)} \tag{2}$$

Assume the observations are independent, the log likelihood for N observations is

$$\begin{aligned} l_N(\theta) &= log(\Pi_{i=1}^{N} P(Female_i|\mathbf{w}_i)) \\ &= \sum_{i=1}^{N} [Female_i * (\theta_0 + \mathbf{w}_i'\theta) - log(1 + exp(\theta_0 + \mathbf{w}_i'\theta))] \end{aligned} \tag{3}$$

I estimate θ on the counts for words through the following objective function[1]:

$$\begin{aligned} \hat{\theta}_\lambda &= argmin_\theta \, (-l_N(\theta)) + \lambda\|\theta\|_1 \\ &= argmin_\theta \, \Sigma_i [log(1 + exp(\theta_0 + \mathbf{w}_i'\theta)) - Female_i(\theta_0 + \mathbf{w}_i'\theta)] + \lambda\|\theta\|_1 \end{aligned} \tag{4}$$

where $\|\theta\|_1 = \sum_{j \geq 1} |\theta^j|$.

Given a word k, we have

$$\frac{\partial P(Female_i = 1|\mathbf{w}_i)}{\partial w_i^k} = P(Female_i = 1|\mathbf{w}_i) * P(Female_i = 0|\mathbf{w}_i) * \theta_\lambda^k \tag{5}$$

where θ_λ^k is the coefficient on w_i^k - the count for word k in post i. Therefore, I estimate the average marginal effect of word k by

$$\frac{1}{N} \Sigma_i P(Female_i = 1|\mathbf{w}_i) * P(Female_i = 0|\mathbf{w}_i) * \hat{\theta}_\lambda^k \tag{6}$$

[1] See Hastie T, Tibshirani R, Friedman J. 2009. *The Elements of Statistical Learning.* Springer. Second Edition. for a detailed discussion of penalized logistic regressions.

Model II. Lasso-regularized Linear Probability Model

Using the same notations as above, I estimate an regularized linear probability model as follows:

$$\hat{\beta}_\lambda = argmin_\beta\, \Sigma_i(Female_i - \beta_0 - \mathbf{w}_i'\beta)^2 + \lambda\|\beta\|_1 \tag{7}$$

where $\|\beta\|_1 = \sum_{j\geq 1} |\beta^j|$.

And the marginal effect of word k on the probability that a post is *Female* is estimated by $\hat{\beta}_\lambda^k$, the coefficient on the regressor w_i^k.

APPENDIX FIGURE 1: Selection of Optimal P-score Cutoff by Mean Squared Error (Lasso-Logistic model on gendered posts identified by the comprehensive list of classifiers.)

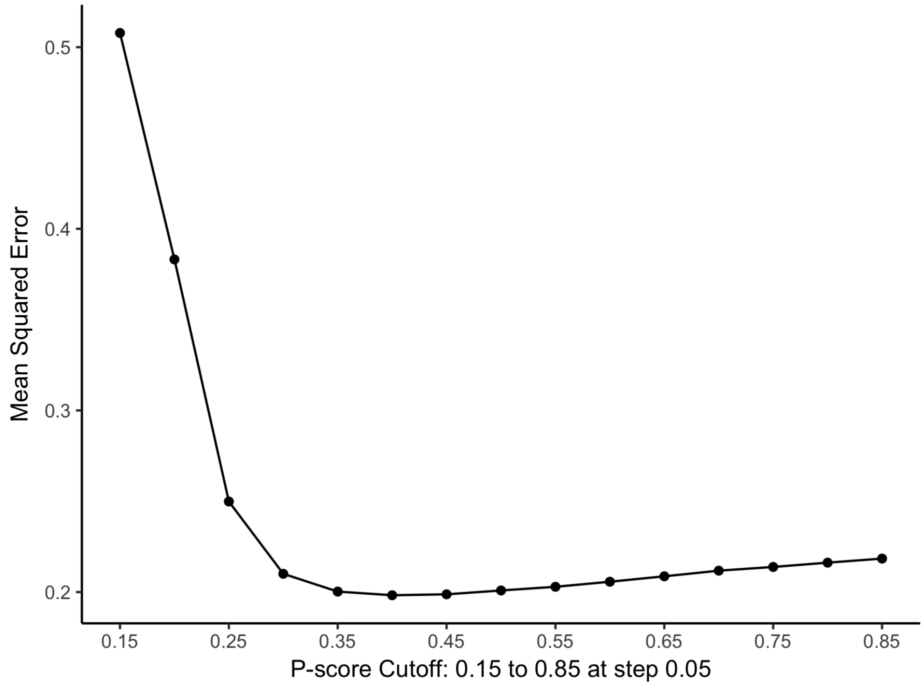

Note: This figure shows the mean squared error (MSE) for predicting gender on the test set of 99,941 gendered posts (a left-out 25% sample) that include only female or only male classifiers from the comprehensive list, at each p-score threshold for assigning a post to *Female* that range from 0.15 to 0.85 with a step size of 0.05. The MSE is minimized at $p = 0.40$. Therefore, I use 0.40 as the threshold to assign genders for 44,081 posts that include both female and male classifiers in the comprehensive list. As a result, 14,028 (31.82%) posts are re-classified to *Female*, and the rest to *Male*.

APPENDIX FIGURE 2: Word Selection by Lasso-Logistic vs. Lasso-Linear (Pronoun Sample)

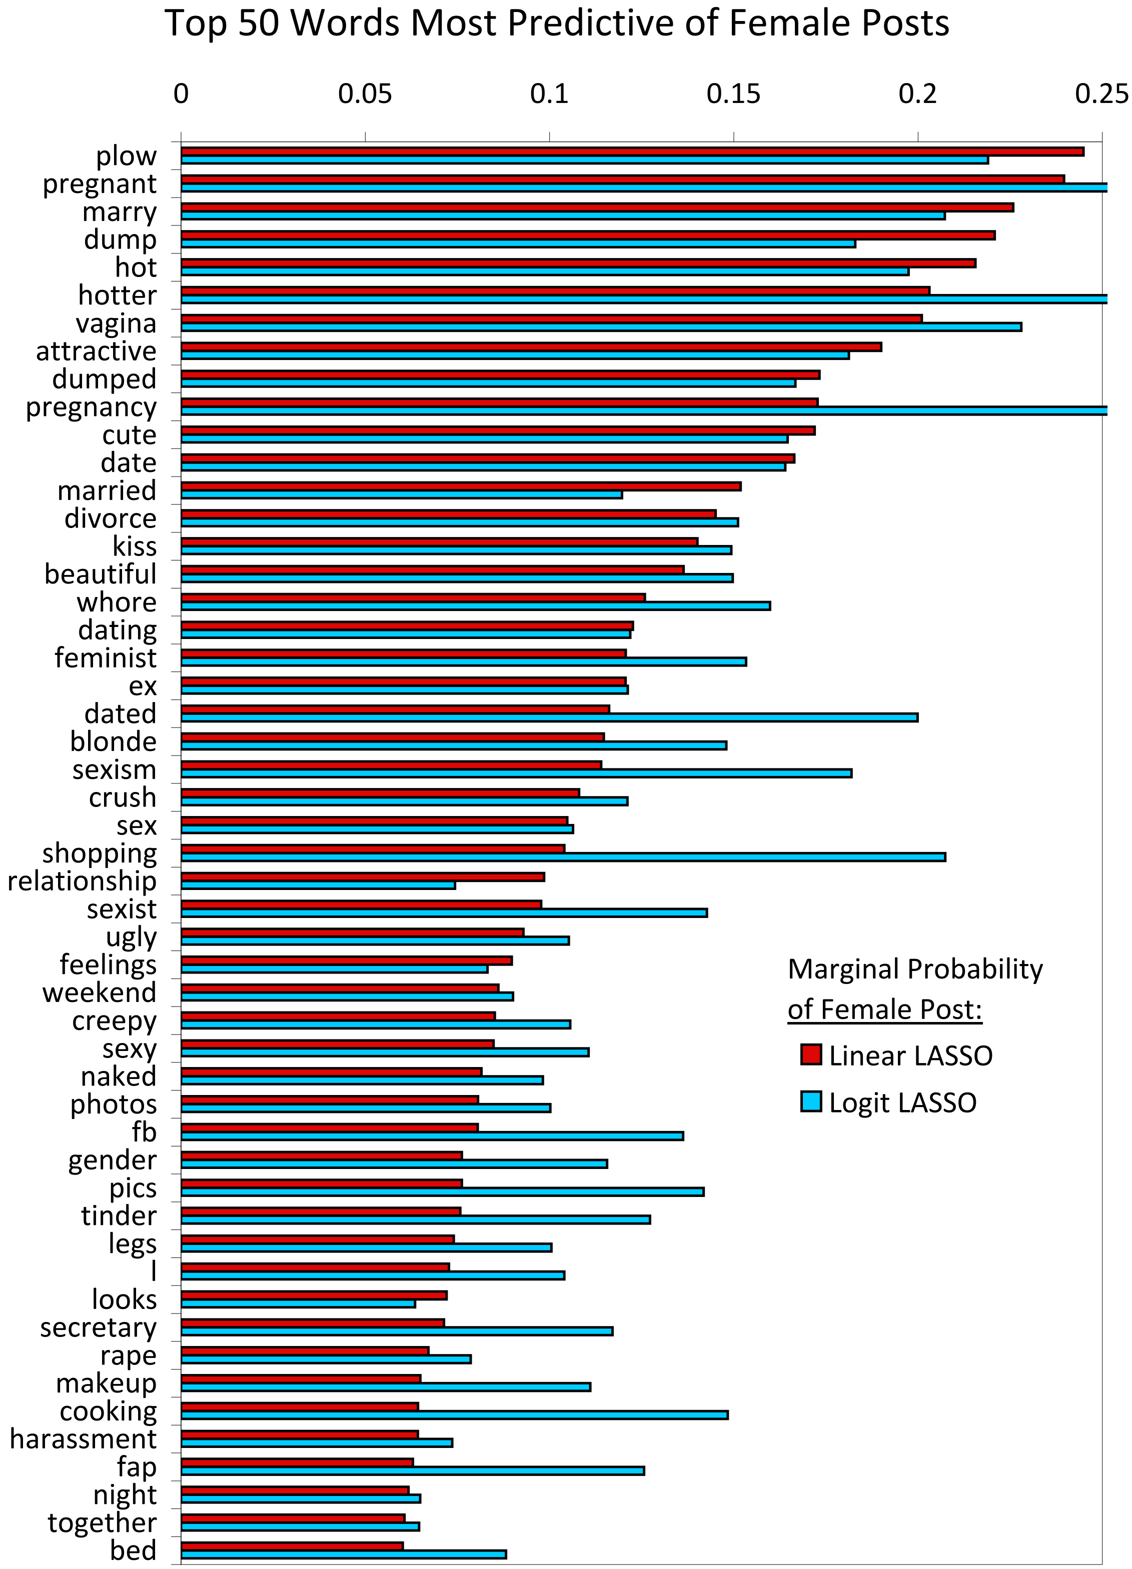

Notes: Each model was trained on 35,850 Female posts and 103,449 Male posts identified by gender pronouns (pronoun sample). The top 50 words above are sorted by the marginal effect of each word estimated by the Linear LASSO model.

APPENDIX FIGURE 3: Word Selection by Lasso-Logistic vs. Lasso-Linear (Pronoun Sample)

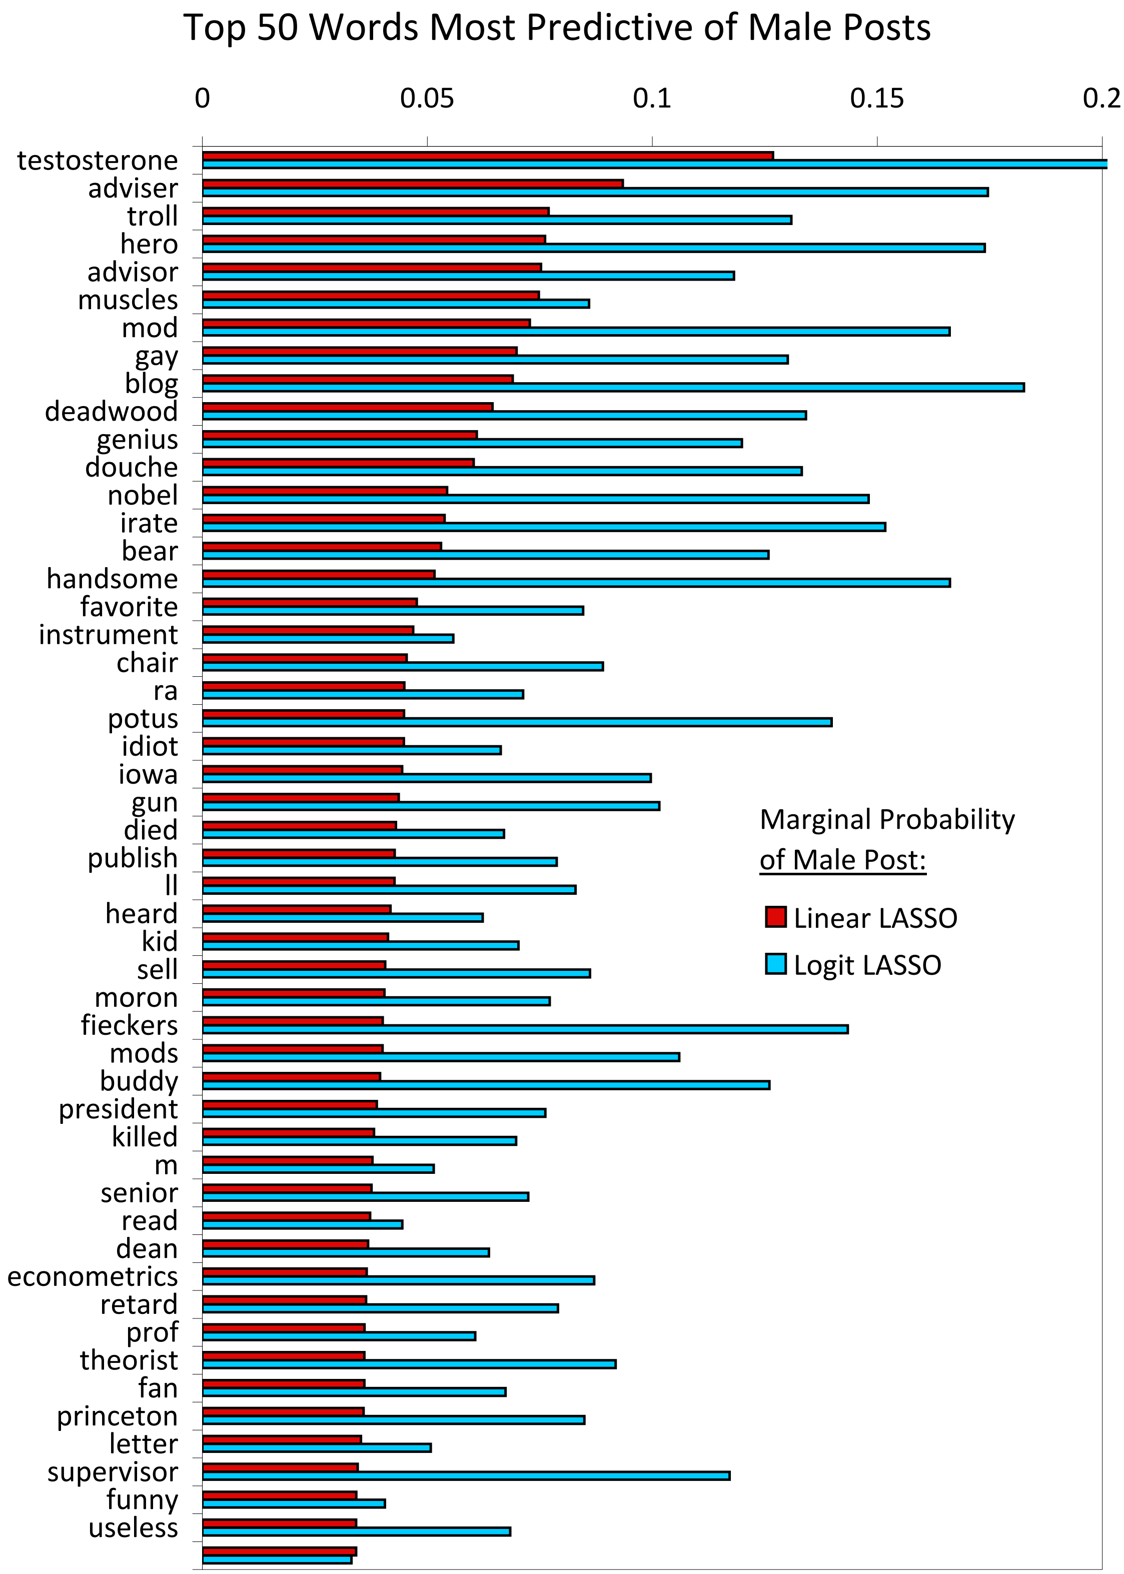

Notes: Each model was trained on 35,850 Female posts and 103,449 Male posts identified by gender pronouns (pronoun sample). The top 50 words above are sorted by the marginal effect of each word estimated by the Linear LASSO model.

APPENDIX TABLE 1: Top 50 *female*(*male*) Words Selected by Lasso-Logistic (Gendered posts are identified by the comprehensive list of classifiers)

Most *female*		Most *male*	
Word	Marginal Effect	Word	Marginal Effect
hotter	0.422	homo	-0.303
pregnant	0.323	testosterone	-0.195
plow	0.277	chapters	-0.189
marry	0.275	satisfaction	-0.187
hot	0.271	fieckers	-0.181
marrying	0.260	macroeconomics	-0.180
pregnancy	0.254	cuny	-0.180
attractive	0.245	thrust	-0.169
beautiful	0.240	nk	-0.165
breast	0.227	macro	-0.163
dumped	0.225	fenance	-0.162
kissed	0.224	founding	-0.160
misogynistic	0.222	blog	-0.157
feminist	0.218	mountains	-0.156
sexism	0.210	grown	-0.156
dated	0.209	frat	-0.155
whore	0.208	handsome	-0.154
sexy	0.202	nba	-0.151
raped	0.200	lyrics	-0.151
attracted	0.198	ferguson	-0.150
slept	0.195	wasn	-0.147
blonde	0.193	supervisor	-0.146
unattractive	0.193	rfs	-0.145
gorgeous	0.192	adviser	-0.141
assaulted	0.191	minnesota	-0.140
cute	0.185	hero	-0.136
vagina	0.184	gay	-0.135
date	0.181	puerto	-0.134
dating	0.181	nobel	-0.129
ugly	0.181	keynesian	-0.128
naked	0.181	sincerely	-0.126
classified	0.179	bashing	-0.126
workforce	0.175	thanks	-0.123
banging	0.175	fiekers	-0.121
impress	0.169	homosexual	-0.121
beauty	0.169	bowl	-0.121
divorce	0.164	nordic	-0.119
feminism	0.164	disability	-0.119
crush	0.163	advised	-0.119
teenage	0.162	inflation	-0.118
dig	0.161	gray	-0.117
sexist	0.160	depth	-0.117
makeup	0.159	wolf	-0.117
cleaning	0.155	curry	-0.116
dump	0.155	teenagers	-0.116
victoria	0.150	wash	-0.116
instagram	0.150	genius	-0.116
tinder	0.149	argues	-0.114
fiecking	0.149	coase	-0.113
shopping	0.149	rip	-0.113

Notes: The top 50 *female* (*male*) words are sorted in descending (ascending) order of their marginal effect - the increase in the probability that the subject of a post is *Female* given an additional occurrence of each word. The model was trained on gendered posts identified by the comprehensive list of gender classifiers.

APPENDIX TABLE 2: Number of Posts that Contain Each of the Top 50 *female*(*male*) Words Selected by Lasso-Logistic (Gendered posts are identified by the comprehensive list of classifiers)

Most *female*			Most *male*		
Word	No. *Female*	No. *Male*	Word	No. *Female*	No. *Male*
hotter	307	31	homo	48	715
pregnant	564	120	testosterone	51	102
plow	274	83	chapters	9	361
marry	1,287	258	satisfaction	59	145
hot	3,613	1,053	fieckers	49	604
marrying	262	49	macroeconomics	19	850
pregnancy	202	61	cuny	8	248
attractive	1,578	417	thrust	6	47
beautiful	1,419	610	nk	3	260
breast	134	48	macro	178	4,282
dumped	361	100	fenance	46	640
kissed	218	50	founding	6	186
misogynistic	66	48	blog	109	1,839
feminist	422	234	mountains	14	90
sexism	269	171	grown	69	394
dated	362	148	frat	59	290
whore	239	148	handsome	103	323
sexy	430	207	nba	16	301
raped	297	155	lyrics	17	111
attracted	415	182	ferguson	10	221
slept	368	85	wasn	32	171
blonde	292	79	supervisor	40	273
unattractive	172	32	rfs	7	284
gorgeous	213	78	adviser	78	712
assaulted	98	52	minnesota	35	703
cute	912	488	hero	47	579
vagina	199	68	gay	406	1,755
date	1,729	835	puerto	7	101
dating	1,423	399	nobel	204	3,379
ugly	1,046	404	keynesian	8	567
naked	376	213	sincerely	55	520
classified	47	96	bashing	15	199
workforce	78	92	thanks	655	4,999
banging	306	109	fiekers	44	406
impress	160	164	homosexual	33	169
beauty	330	193	bowl	24	203
divorce	673	192	nordic	103	537
feminism	264	127	disability	27	117
crush	320	207	advised	33	227
teenage	168	116	inflation	41	1,000
dig	152	176	gray	34	108
sexist	469	358	depth	17	257
makeup	174	66	wolf	19	144
cleaning	175	169	curry	12	143
dump	503	339	teenagers	36	113
victoria	40	49	wash	74	204
instagram	100	63	genius	92	1,007
tinder	301	110	argues	23	313
fiecking	377	226	coase	7	200
shopping	165	129	rip	56	484

Notes: This table shows the number of *Female* posts and the number of *Male* posts that contain each of the top 50 *female* or *male* terms selected by Lasso, in the same order as in Appendix Table 1. Using the comprehensive list of gender classifiers, I identified 103,584 *Female* posts and 341,226 *Male* posts.

APPENDIX TABLE 3: Most Frequent Words in *Female* (*Male*) posts, identified by the comprehensive list of classifiers

Most common in *Female*			Most common in *Male*		
Word	No. *Female*	No. *Male*	Word	No. *Female*	No. *Male*
life	4,034	7,644	work	3,800	13,989
work	3,800	13,989	paper	1,503	11,727
hot	3,613	1,053	job	3,091	10,313
love	3,297	4,274	economics	1,120	9,808
sex	3,103	1,535	great	2,323	9,181
job	3,091	10,313	best	2,558	8,552
feel	2,574	5,167	research	1,407	8,238
best	2,558	8,552	school	2,446	8,228
school	2,446	8,228	market	1,750	7,954
kids	2,441	2,200	life	4,034	7,644
great	2,323	9,181	phd	1,751	7,295
married	2,231	1,207	papers	854	7,177
friends	2,048	2,504	econ	1,133	6,950
nice	1,978	4,590	students	1,474	6,889
money	1,951	6,011	theory	415	6,347
home	1,778	2,734	money	1,951	6,011
phd	1,751	7,295	data	729	5,648
market	1,750	7,954	student	1,560	5,607
date	1,729	835	economist	855	5,539
family	1,653	2,685	wrong	1,344	5,487
attractive	1,578	417	economists	697	5,461
student	1,560	5,607	course	1,320	5,416
relationship	1,506	1,169	question	1,109	5,257
paper	1,503	11,727	idea	1,158	5,184
students	1,474	6,889	feel	2,574	5,167
happy	1,452	2,536	economic	466	5,152
dating	1,423	399	department	935	4,985
beautiful	1,419	610	university	955	4,970
friend	1,412	2,423	r	682	4,774
research	1,407	8,238	nice	1,978	4,590
single	1,373	2,578	finance	357	4,469
wrong	1,344	5,487	working	1,282	4,465
children	1,337	1,449	field	547	4,339
course	1,320	5,416	policy	504	4,330
young	1,315	2,751	macro	178	4,282
marry	1,287	258	love	3,297	4,274
working	1,282	4,465	model	463	4,210
social	1,257	3,590	tenure	930	3,891
fat	1,237	1,170	public	820	3,877
aspie	1,235	1,412	journal	324	3,787
idea	1,158	5,184	professor	679	3,781
marriage	1,150	614	class	1,115	3,614
age	1,142	1,881	social	1,257	3,590
econ	1,133	6,950	harvard	418	3,533
economics	1,120	9,808	business	546	3,478
class	1,115	3,614	math	394	3,421
question	1,109	5,257	offer	777	3,401
college	1,095	2,651	nobel	204	3,379
ugly	1,046	404	able	979	3,320
experience	1,043	2,876	academic	654	3,280

Notes: The words that are most common in *Female* (*Male*) are sorted by the number of *Female* (*Male*) posts they appear in. Using the comprehensive list of gender classifiers, I identified 103,584 *Female* posts and 341,226 *Male* posts.

APPENDIX TABLE 4: Top 50 *female*(*male*) Words Selected by Lasso-Logistic (Gendered posts are identified by pronouns only)

Most *female*		Most *male*	
Word	Marginal Effect	Word	Marginal Effect
pregnancy	0.292	knocking	-0.329
hotter	0.289	testosterone	-0.204
pregnant	0.258	blog	-0.183
hp	0.238	hateukbro	-0.176
vagina	0.228	adviser	-0.175
breast	0.220	hero	-0.174
plow	0.219	cuny	-0.173
shopping	0.207	handsome	-0.166
marry	0.207	mod	-0.166
gorgeous	0.201	homo	-0.160
dated	0.200	rfs	-0.154
marrying	0.198	irate	-0.152
hot	0.197	nobel	-0.148
dump	0.183	dictator	-0.144
sexism	0.182	fieckers	-0.143
attractive	0.181	spell	-0.143
sperm	0.171	potus	-0.140
dumped	0.167	nk	-0.137
intimate	0.167	repec	-0.137
cute	0.165	minnesota	-0.135
date	0.164	advising	-0.135
whore	0.160	deadwood	-0.134
commonly	0.159	ego	-0.133
commodities	0.159	douche	-0.133
consent	0.153	punch	-0.131
feminist	0.153	troll	-0.131
classified	0.152	gay	-0.130
divorce	0.151	gays	-0.129
beautiful	0.150	beard	-0.127
kiss	0.149	writings	-0.127
victoria	0.149	blanket	-0.127
cooking	0.148	bowl	-0.127
blonde	0.148	buddy	-0.126
yoga	0.147	bear	-0.126
oct	0.144	ferguson	-0.125
sexist	0.143	legend	-0.124
pics	0.142	assumes	-0.123
university's	0.140	westerners	-0.123
improvements	0.140	rip	-0.121
fb	0.136	sins	-0.120
aej	0.136	genius	-0.120
yahoo	0.134	evolution	-0.119
cum	0.133	advisor	-0.118
rct	0.133	supervisor	-0.117
activist	0.133	calculus	-0.117
flirting	0.132	goals	-0.116
feminism	0.129	decency	-0.116
tinder	0.127	penalty	-0.116
flowers	0.126	injured	-0.113
instagram	0.126	depth	-0.113

Notes: The top 50 *female* (*male*) words are sorted in descending (ascending) order of their marginal effect - the increase in the probability that the subject of a post is *Female* given an additional occurrence of each word. The model was trained on gendered posts identified by feminine or masculine pronouns only.

APPENDIX TABLE 5: Number of Posts that Contain Each of the Top 50 *female*(*male*) Words Selected by Lasso-Logistic (Gendered posts are identified by pronouns only)

Most *female*			Most *male*		
Word	No. *Female*	No. *Male*	Word	No. *Female*	No. *Male*
pregnancy	106	27	knocking	6	82
hotter	120	31	testosterone	15	31
pregnant	270	98	blog	89	1,244
hp	26	14	hateukbro	0	70
vagina	137	41	adviser	66	591
breast	62	30	hero	32	412
plow	146	60	cuny	3	104
shopping	99	69	handsome	41	170
marry	557	191	mod	30	384
gorgeous	110	41	homo	31	162
dated	194	86	rfs	5	137
marrying	117	34	irate	24	235
hot	1,309	658	nobel	125	1,944
dump	369	215	dictator	6	167
sexism	87	76	fieckers	20	201
attractive	547	246	spell	26	127
sperm	46	22	potus	20	202
dumped	240	88	nk	0	119
intimate	49	27	repec	8	176
cute	463	298	minnesota	21	282
date	902	477	advising	10	193
whore	123	112	deadwood	38	426
commonly	25	57	ego	47	245
commodities	12	28	douche	48	288
consent	83	62	punch	25	153
feminist	162	128	troll	206	1,606
classified	33	56	gay	163	737
divorce	376	147	gays	5	78
beautiful	524	346	beard	14	99
kiss	308	148	writings	3	157
victoria	18	17	blanket	9	55
cooking	72	44	bowl	14	104
blonde	155	51	buddy	50	193
yoga	53	38	bear	104	736
oct	23	143	ferguson	10	126
sexist	145	161	legend	13	117
pics	121	77	assumes	12	139
university's	30	68	westerners	5	39
improvements	14	43	rip	33	218
fb	120	84	sins	5	88
aej	34	82	genius	50	650
yahoo	23	42	evolution	15	152
cum	67	60	advisor	286	2,145
rct	15	26	supervisor	34	199
activist	41	89	calculus	10	246
flirting	103	27	goals	38	304
feminism	68	56	decency	5	64
tinder	106	34	penalty	14	170
flowers	73	42	injured	9	158
instagram	65	39	depth	11	143

Notes: This table shows the number of *Female* posts and the number of *Male* posts that contain each of the top 50 *female* or *male* terms selected by Lasso, in the same order as in Appendix Table 4. Using gender pronouns, I identified 49,993 *Female* posts and 145,382 *Male* posts.

APPENDIX TABLE 6: Most Frequent Words in *Female* (*Male*) posts, identified by pronouns only

Most common in *Female*			Most common in *Male*		
Word	No. *Female*	No. *Male*	Word	No. *Female*	No. *Male*
work	2,227	8,018	work	2,227	8,018
life	2,017	4,133	paper	1,030	6,500
love	1,762	2,055	job	1,609	5,517
job	1,609	5,517	great	1,371	4,840
feel	1,523	2,339	economics	640	4,696
sex	1,377	831	best	1,320	4,423
great	1,371	4,840	school	1,334	4,314
school	1,334	4,314	research	828	4,270
best	1,320	4,423	papers	592	4,194
hot	1,309	658	life	2,017	4,133
married	1,116	678	students	766	3,867
student	1,109	3,781	phd	968	3,837
friends	1,088	1,459	student	1,109	3,781
nice	1,055	2,412	market	702	3,706
paper	1,030	6,500	economist	542	3,345
kids	1,009	1,236	money	975	3,307
home	989	1,562	course	769	3,146
money	975	3,307	wrong	827	3,144
friend	974	1,951	idea	702	3,009
phd	968	3,837	department	620	2,926
date	902	477	econ	587	2,820
relationship	880	644	theory	258	2,787
family	863	1,601	question	619	2,717
happy	850	1,334	professor	485	2,578
research	828	4,270	university	640	2,533
wrong	827	3,144	economists	338	2,482
course	769	3,146	tenure	618	2,462
students	766	3,867	working	702	2,449
market	702	3,706	nice	1,055	2,412
working	702	2,449	economic	257	2,376
idea	702	3,009	feel	1,523	2,339
economics	640	4,696	data	341	2,278
university	640	2,533	field	324	2,225
department	620	2,926	advisor	286	2,145
question	619	2,717	class	606	2,106
tenure	618	2,462	offer	522	2,097
class	606	2,106	public	488	2,077
couple	598	1,300	policy	309	2,069
papers	592	4,194	love	1,762	2,055
econ	587	2,820	journal	219	1,987
mind	583	1,607	friend	974	1,951
marriage	580	368	able	542	1,950
dating	573	242	nobel	125	1,944
marry	557	191	r	363	1,933
young	547	1,498	published	282	1,930
attractive	547	246	smart	535	1,904
economist	542	3,345	editor	201	1,837
able	542	1,950	stupid	456	1,822
social	538	1,760	academic	378	1,801
smart	535	1,904	social	538	1,760

Notes: The words that are Most common in *Female* (*Male*) are sorted by the number of *Female* (*Male*) posts they appear in. Using gender pronouns, I identified 49,993 *Female* posts and 145,382 *Male* posts.