The Welfare Effects of Gender-Inclusive Innovation: Evidence from Patents and Academic Publications

Matt Marx
Cornell University & NBER

Frank Müller-Langer
University of the Bundeswehr Munich & MPI-IC

Joel Waldfogel
University of Minnesota, NBER, & ZEW

July 14, 2026

PRELIMINARY DRAFT

Abstract

Since 1970 the share of mostly-female-inventor patents has risen from five to over nine percent, while the mostly-female share of academic articles has risen from under five to 20 percent. We explore the effects of these influxes. First, female innovations are cited nearly as much as male innovations, and particularly by female innovators; additional innovations "expand the market" by attracting citations, rather than simply diverting citations from existing innovations; the quality of innovations is unpredictable at entry; and direct evidence that growing female innovation displaces male innovation is absent. Second, we develop an equilibrium model of innovation demand and entry. We remove female innovations to bring the relative female contribution to its pre-1970 level, and we simulate the male entry response with a model of entry order based on quality predictability. The female influx into patenting raises innovator surplus by 0.82 percent overall and by similar amounts by citer gender, with particularly high benefits (over 3.5 percent) in pharmaceuticals and basic materials chemistry. The female influx into academic authorship raises overall surplus by 2.3 percent and by 3.4 and 2.0 percent, respectively, for female and male academic authors, with particularly large benefits in academic fields related to law and education. Finally, the marginal benefit of female innovation is large in some areas – and some STEM areas – with still-low female participation.

Introduction

Innovation is essential to growth, so it would be valuable to attract additional qualified people to participation in scientific and innovative activity. While women have become roughly half of the labor force, they have been slower to move into technical, innovative, and creative occupations. Female authors accounted for under five percent of the academic articles published in 1965, with much lower shares in STEM fields. Half a century later—in 2015—female-authored articles accounted for 20 percent of new articles. The growth in female-inventor entry into patenting is similar but less pronounced. The share of patents with mostly-female inventors rose from 5 percent in 1976 to over 9 percent in 2020. There is reason to believe that the low female participation has caused valuable innovations to have been forgone (Bell et al., 2019). The goal of this paper is to use the growth in female participation in patenting and academic authorship – and the resulting impacts on the pool of extant articles and patents (collectively "innovations") – to measure the benefit of more inclusive scientific innovation.

We examine the effect of the female influx into innovation in two parts. First, we provide extensive descriptive evidence on both the growth in female innovative activity and its apparent value. Second, we quantify the relative value of the female influx into authorship with an equilibrium model of demand for citations and entry into innovation. We treat the citing behavior of new inventors and authors as demand, and we treat the extant patents and articles as the choice set. We compare the "innovator surplus" (IS) that citing innovators derive from the contemporary choice set compared with a counterfactual without the female influx. The counterfactual environment is less valuable, the more that valuable female innovations are removed; but simple removal of female innovations will overstate the benefit of the female influx if male entry would respond, as the removal of the female influx raises the expected return to innovation. Hence, the equilibrium counterfactual exercise requires two distinct

elements, a removal of female innovations to bring each female vintage to its value relative to male innovation prior to 1970 and a model of equilibrating male entry based on the expected citations to marginal male innovations in relation to entry costs. Male innovations enter in order of predicted quality, and the expected citations to the marginal male innovation is a high-dimensional random variable. We develop a tractable approach to estimating expected citations to marginal entrants for entry levels beyond the status quo, which allows us to model the value of counterfactual male entry and therefore the equilibrium welfare effect of the female influx.

We undertake these exercises using two broad datasets. For academic articles, we use the Microsoft Academic Graph (MAG) on the citations received by all 65.2 million journal articles published between 1965 and 2015 in 40 distinct academic areas. We classify each article according to its authors' genders, and we know how many of the citations it received each year came from male vs female-authored citing papers. We use forename-country-based algorithms and an LLM to classify the genders of 66.7 million authors. For patents we use the gender-identified data on 7 million patents granted between 1976 and 2020 created by the US Patent and Trademark Office.

We have six descriptive findings suggesting that the female influx is valuable. First, there is a substantial female influx into almost all areas of innovative activity in the last half century. Second, innovations from women attract nearly as many citations, per innovation, as those by men although women cite work by women substantially more. Third – and related – increases in the female-authored share of innovations, by field and vintage, have been accompanied by roughly proportional increases in the share of citations garnered by female innovations, which is consistent with the female influx delivering high-quality research. Fourth, fields and years with more extant innovations "expand the market" by attracting more citations, which is consistent with the additional innovations being valuable. Fifth, the quality of innovation is highly unpredictable at the time of entry. Sixth, we see no evidence

of displacement of male innovation by the female influx. Using the equilibrium model, we estimate that the female influx into patenting has raised patent innovator surplus by 0.8 percent overall, with similar effects for male and female inventors. The female influx into academic research has raised overall author IS by 2.29 percent, while raising female and male IS by 3.39 and 2.02 percent, respectively. The marginal benefit of innovation is high in some areas with still-low female participation, suggesting the benefit of continued influxes.

The paper proceeds in five sections after the introduction. Section 1 provides background, including historical context and literature. Section 2 describes a theoretical setup that implies empirical questions and provides a framework for welfare analysis. Section 3 describes the data used in the study. Section 4 presents descriptive evidence on the female influx and its impact on citations. Section 5 presents estimates of the equilibrium model and the counterfactual exercises.

1 Background

1.1 Historical context: the supply of female innovators

Women's participation in science and engineering was limited in the early and mid-20^{th} century and was instead concentrated in a narrow set of occupations (Costa, 2000; Goldin, 2014). For instance, women accounted for less than a percent of engineers and chemists in 1940 and only 1.2 percent of engineers in 1950 (Pidgeon, 1954; Blank and Stigler, 1957). Despite their absence, women appear to have been a large and untapped pool of scientific potential: Altonji et al. (2016) and Jiang (2021) find that gender differences in ability and ability sorting explain only a small portion of the gender gap in STEM majors, implying that many women with STEM-relevant skills do not select into STEM fields.

Performance on standardized tests for graduate school admission provides additional evidence of underutilized potential (Hyde and Linn, 2006; Hyde et al., 2008). By the early

1980s, women were taking the Graduate Record Exam (GRE) in greater numbers than men; and female test-takers have come to account for a large share of high score recipients. Although men outscore women on average, there is substantial dispersion; and both means and standard deviations are reported by gender for the math and verbal tests in some years (Bleske-Rechek and Browne, 2014). As Figure 1 shows, given the rather substantial dispersion in scores – along with the large entry of women into test-taking – the female influx into test taking roughly doubled the number of US test takers receiving a quantitative score over 700.[1] This score is high enough to gain entry to most top engineering programs.[2]

Since the late 1960s, women's educational attainment patterns have shifted substantially toward more remunerative and career-oriented fields (Mammen and Paxson, 2000; Goldin, 2014). Women have come to account for growing shares of PhDs granted over the past half century, as Figure 2 shows.[3] The female share of PhD grantees varies across fields but has risen in all of them since 1960. Women now account for over 75 percent in psychology, nearly 70 percent in education, and roughly 55 percent in the arts. Among the fields more directly relevant to innovative activity, the female share exceeds 50 percent in health and biology and by 2018 had reached roughly 20-25 percent in math, computer science, physical science, and engineering.[4] It is clear that the cadre of doctorate holders, including those in STEM fields,

[1] Assuming that the math and verbal scores are distributed normally, one can calculate the numbers of men and women with scores above any absolute level. Suppose that G_m and G_f are the numbers of men and women taking the exam, μ_m^v and σ_m^v are the average and standard deviation of the men's verbal score in a year, and Φ is the normal cumulative distribution function. Then the number of women scoring above, say, 700 on the verbal component in a year is given by

$$G_f \left(1 - \Phi(\frac{700 - \mu_f^v}{\sigma_g^v})\right),$$

and the numbers of men or women scoring above any particular threshold in a year is calculated analogously.

[2] According to `https://magoosh.com/gre/gre-scores-for-engineering-programs/`, the average GRE scores for most kinds of engineering average between 152 and 161 on the new scale. This maps to 670-770 on the old scale. See `https://www.prepscholar.com/gre/blog/old-gre-to-new-gre-score-conversion-charts/`.

[3] In 1960, 6 of 959 US PhDs in engineering were awarded to women. In 1949, women earned 175 of 52,246 engineering bachelor's degrees, 15 of 4,496 engineering master's degrees, and just 1 of 417 engineering doctorates (Pidgeon, 1954).

[4] See Digest of Education Statistics, Tables 325.35 (computer science), 325.45 (engineering), 325.65 (math-

is growing more female over time. This raises the possibility of harnessing the innovative potential of a previously less-engaged group.

1.2 Relevant literature

A growing literature has documented both the gaps and growth in women's participation in the scientific enterprise (Huang et al., 2020). Larivière et al. (2013) document persistent gender disparities in scientific publishing, with women comprising less than 30% of fractionalized authorships on papers from 2008–2012. Women were moreover underrepresented as first authors, and received fewer citations overall when occupying key authorship positions (first, last (i.e., PI), or solo). Shen (2013) quantified ongoing inequalities in U.S. academia, highlighting that despite gains in entry-level representation, women remain underrepresented in senior faculty positions, suggesting a "leaky pipeline" despite efforts to close the gap. Related work explored underlying reasons for the gap. National Academy of Sciences (2006) concluded that barriers arise largely from implicit bias and legacy institutional structure as opposed to intentional discrimination or innate ability differences. Ceci and Williams (2011) for example claimed that gender imbalance in math-intensive fields stems not primarily from overt discrimination in peer review, grants, or hiring; rather, individual choices (e.g., regarding time allocation to family). Kim and Moser (2021) document the role of children in women's scientific careers.

Similar patterns obtain not just in the publishing of scientific discoveries but their subsequent commercialization. Several studies suggest that women patent their scientific discoveries at lower rates (Thursby and Thursby, 2005; Whittington and Smith-Doerr, 2005; Ding et al., 2006; Fechner and Shapanka, 2018). Given that the commercialization of science is not limited to patenting, Tartari and Salter (2015) report less engagement in consulting and other translational activities. Fieldwork establishes that gender differences in institutional

ematics), and 325.70 (physical sciences), `https://nces.ed.gov/programs/digest/`.

and individual characteristics favoring the commercialization of academic research may contribute to this gender gap (Murray and Graham, 2007). Koffi and Marx (2025) report a double-digit gender gap in the commercialization of scientific discoveries as captured via patent-paper "pairs" (Ducor, 2000).

At the same time, the gender gap in scientific production seems to be shrinking, albeit not according to every metric. Huang et al. (2020) note that even as women's annualized participation in science increased substantially, gender gaps in total career productivity and impact paradoxically widened (e.g., from a ∼10% productivity gap in the 1950s to 35% favoring men in the 2000s). Iaria et al. (2022) notes similarly that only one percent of academics were women in 1900, and the share of women increased slowly until 1969 but faster in the last three decades of the 20th century. Still, by the end of the 20th century women remain less prevalent in the sciences relative to men, particularly in prestigious universities.

Despite the substantial progress to date in characterizing the magnitude and possible mechanisms underlying the gender gap in scientific production, less work addresses welfare effects. Some work explores the impact of female participation on the direction of innovative activity (Einiö et al., 2019; Koning et al., 2021). Existing work does not quantify the benefit of female participation in innovative activity.

2 Theory

2.1 Analogy to demand and entry

We analyze the value of the female influx using an analogy to the IO approach to studying product markets. We view innovators' citations to prior innovations as demand, and we treat the decision to create an innovation as entry. In the following discussion we refer to articles and patents collectively as "innovations," and we refer to authors and inventors as

"creators" or "innovators." We label an innovation as *female* (f) if its creator team is mostly female and *male* (m) otherwise. Thus, m is residual, i.e., mostly-male, gender-balanced, or gender-unidentified teams. It should be read as "not mostly-female" rather than all-male.

Demand. The creators of each new innovation are making a series of choices of which of the already-existing innovations to cite. The innovations $j \in J$ available to be cited at time t are the supply side of the market, or the product "choice set." In order to create a tractable entry model, we employ a nested logit approach to demand, as in Berry (1994). For each of the instances in which a creator i might make a citation to innovation j, there is a utility of doing so:

$$u_{ij} = \delta_j + \zeta_{ij} + (1-\sigma)\epsilon_{ij},$$

where $\delta_j = x_j\beta - \alpha p_j + \xi_j$ is the mean utility of innovation j, ζ_{ij} is the common component of utility across innovations, and the parameter σ quantifies substitutability. When $\sigma = 0$, this is ordinary logit; as $\sigma \to 1$, innovations become more substitutable for one another. The utility of the outside good, not citing an innovation, is $u_{i0} = 0$. Each innovation j has a vintage in which it appeared (v) as well as a topic area (a). While there is no explicit price of citing an innovation, there is an implicit price, as innovations cite a finite number of innovations. The lack of a literal price prevents us from estimating the money metric parameter α.

We are interested in the effect of the influx of female-created innovations on the benefit that innovators derive from the available choice set. We make an analogy between consumer surplus and the "innovator surplus" that citing innovators derive from the available choice set. Innovator surplus (IS), by analogy to CS, is proportional to

$$IS(D) \propto \frac{M}{\alpha} \ln(1 + D^{1-\sigma}),$$

where D is an importance-weighted sum of the extant innovations available at some point in calendar time, and M is market size. In particular, $D = \Sigma_{j \in J} e^{\delta_j/(1-\sigma)}$, and J is the set of extant articles at the time that citation decisions are being made. Note that D has subcomponents by field (a), vintage (v), and gender, e.g., D^f_{va}, such that, for example, $D_a = \Sigma_{g \in f,m} \Sigma_v D^g_{va}$.

2.2 Supply/entry

We are interested in the benefit that ensues from the influx of women into innovative activity. To this end, we compare contemporary environments – in particular, contemporary innovation choice sets that include a female influx – with counterfactual environments in which the influx did not occur. In addition to a model of demand that allows welfare evaluation of choice sets, this exercise requires the answers to two distinct questions. The first is which of the extant female innovations to remove to undo the female influx. The second is the modeling of the male response to the removal of the female influx. These removal and entry margins are quite different conceptually, and they require different analytical approaches.[5]

2.2.1 Extensive margin: which female innovations to remove

Each vintage v and topic area a delivers N^f_{va} female and N^m_{va} male innovations. Our counterfactual seeks to remove female innovations from the status quo to mimic, say, a pre-influx period v, which we term v_0. If female innovations were equally good relative to male innovations over time, we could simply remove female innovations from each post-v_0 vintage at random. Then, for example, we could replace observed N^f_{av} with $\tilde{N}^f_{av} = N^m_{av} \left(\frac{N^f_{av_0}}{N^m_{av_0}} \right)$. Yet, it is unnecessary to make such an assumption; we can instead take a direct empirical approach that relates the number of innovations to their value.

[5]The model builds on the approaches in Aguiar and Waldfogel (2018) and Waldfogel (2025).

The term δ_j is the realized quality of innovation j.[6] Define δ_j^f and δ_j^m as qualities of particular female and male innovations. We can therefore calculate the total importance-weighted values of D_{av}^f and D_{av}^m for each vintage. We can implement the no-female-influx counterfactual by replacing D_{av}^f with $D_{1av}^f = D_{av}^m \left(\frac{D_{av_0}^f}{D_{av_0}^m}\right)$. This approach removes the value of female innovation in post-v_0 vintages to bring the relative value of female innovations in each area and vintage in line with the ratio for period v_0. Then the no-female-influx value of D is described by $D_1^f = \Sigma_a \Sigma_v D_{1av}^f$. Note that we need not determine which of the female innovations to remove; rather, we need only determine the aggregate value of the female innovations to remove from the status quo choice sets.

The removal of female innovations would, by itself, lead to a naive calculation of the welfare effect of the female influx. Denoting status quo by 0 and the no-female-influx counterfactual by 1, the percent change in innovator surplus from the elimination of the female influx alone would be given by

$$\%\Delta IS = 100 \times \left[\frac{\ln(1 + (D_0^f + D_0^m)^{1-\sigma})}{\ln(1 + (D_1^f + D_0^m)^{1-\sigma})} - 1\right]. \tag{1}$$

Because $D_1^f < D_0^f$ the no-male-entry-response $\%\Delta IS$ must be positive. Yet, the reduction in D would raise the marginal return to male entry, so a comparison of the status quo to a counterfactual environment with simple removal of the female influx does not provide a satisfactory equilibrium characterization of the benefit of the female influx. That instead requires a model of additional male entry.

2.2.2 Intensive margin: additional male entry

Innovators launch another innovation if the return they expect exceeds the cost. Define the cost of entry as K. Returns to entry come in the form of citations c. For simplicity, suppress

[6] We will observe these "product qualities" δ_j as $\delta_j = \ln(s_j) - \sigma \ln(s_j/1-s_0) - \ln(s_0)$, where $s_j = {}^{\text{citations to j}}/_{M}$, σ is a nested logit substitution parameter, and M is market size.

vintage and field subscripts and define K_0^f and K_0^m as the gender-specific entry costs in the status quo. Further, define N^f and N^m as the numbers of female and male innovations, respectively; and define $D^m(N^m) = \Sigma_{j=1}^{N^m} e^{\delta_j^m/(1-\sigma)}$ and $D^f(N^f)$ analogously.

When N^m male and N^f female innovations enter, the realized citations to the marginal entering (male) innovation is given, in this model, by

$$q^m(N^m, N^f) = \left[M \frac{e^{\delta_{N^m}/(1-\sigma)}}{D^m(N^m) + D^f(N^f)} \frac{(D^m(N^m) + D^f(N^f))^{1-\sigma}}{1 + (D^m(N^m) + D^f(N^f))^{1-\sigma}} \right], \tag{2}$$

where M is market size. This setup requires an entry order: δ_{N^m} is the quality of the last entering male product. Given an entry order, if the quality of each entering innovation is predictable and known ex ante, Equation (2) is deterministic even prior to entry. In the more realistic case that quality is not fully predictable, the expectation of q becomes a random variable that depends on both the entry order and the quality realizations for all of the entering innovations. Thus, $E[Q^m(.,.)]$ depends on an $(N^f + N^m)$-dimensional integral.

Our major use for this model is a comparison of innovator surplus from the status quo – which includes the female influx – to a counterfactual environment with a) the female influx removed and the female set of innovations held constant, and b) with the male innovation set allowed to expand to fill the void created by the elimination of the female influx. These exercises dictate a few things that a model must provide.

First, in order to re-equilibrate the male entry model, we need to have a measure of the cost threshold governing male entry in the status quo, K_0^m. Second, we need a way to calculate the expected return function for male entry beyond the status quo N_0^m, conditional on female entry being held fixed. That is, we need the function $E[Q^m(N^m; D_1^f)]$ for $N^m > N_0^m$, where D_1^f is the value of the fixed no-influx level of female entry, and N^m is the adjustable male entry.

To calculate the status quo male cost threshold, we define an equilibrium level of male

entry N^m such that

$$E[Q^m(N^m + 1; N_0^f)] < K_0^m < E[Q^m(N^m; N_0^f)]. \tag{3}$$

That is, we assume that potential innovations with expected returns above K_0^m enter, while those with expected returns below that threshold do not enter.[7] The value of entering female products in the status quo is summarized by D_0^f, while the no-female-influx value is D_1^f. Then the realized $E[Q^m(N^m; D_1^f)]$ function is straightforward, once we have the order in which the male products enter.

Predictability and entry order. Entry order depends on the predictability of innovation quality, as two extreme examples make clear. Suppose first that realized quality were completely predictable. Then $\delta_1^m > \delta_2^m > \cdots > \delta_n^m$, and the relationship between the value of entering products and the number of entering innovations depends on entry order. Thus, of course, $D^m(N^m)$ depends on this entry order. Because of the monotonically declining qualities, $D^m(N^m)$ rises quickly but at a rapidly decreasing rate with perfect predictability. As a result, the (deterministic) function $E[Q^m(N^m; .)]$ falls rapidly in N^m; and for a given status quo male entry (N_0^m), implied entry costs K_0^m are low.

At the other extreme, if quality were entirely unpredictable at entry, then there would be no particular entry order. In expectation, marginal entrants would have quality equal to the average quality of inframarginal entrants. Then $D^m(N^m) = N^m e^{\bar{\delta}^m/(1-\sigma)}$, where $\bar{\delta}^m$ is the average quality of male entry. Hence $\frac{\partial D^m}{\partial N^m} > 0$ (and constant), and $\frac{\partial^2 D^m}{\partial N^{m2}} = 0$. In contrast to the fully predictable case, $E[Q^m(N^m; .)]$ would fall slowly; and for a given level of status quo male entry (N_0^m), the entry cost we would infer (K_0^m) is large.

Given the well-known (largely) unpredictability of innovation and product quality (Harhoff et al., 1999; Aguiar and Waldfogel, 2018), the case with partial predictability is more rele-

[7] The term N^m is sufficiently large that the interval is narrow.

vant, albeit more complicated. Suppose that expected quality of innovation j, δ_j, is given by $\delta_j^{'}$, which is a prediction knowable at the time of entry. Then, given the fixed no-influx level of female entry, potential (male) entrants would enter in declining order of $\delta_j^{'}$. Define ρ as the correlation between δ_j and $\delta_j^{'}$. Then the term $D^m(N^m)$ depends on entry order which, in turn, depends on ρ, so that $D^m(N^m) = D^m(\{\delta_j^m\}, \rho)$, where $\{\delta_j^m\}$ is the set of N^m realized innovation qualities, and ρ summarizes predictability.

We seek a function $E[Q^m(N^m; \rho, D^f(N_1^f))]$ showing expected returns for any level of male entry N^m given predictability and the level of female entry. Using this function, along with status quo male entry costs $K_0^m = K^m(N_0^m, D_0^f, \rho)$, we can calculate counterfactual male entry N_1^m under the female influx from $K_0^m = q^m(N_1^m, D^f(N_1^f), \rho)$. We discuss details of implementation in Section 5.

Finally, we calculate welfare, or innovator surplus as

$$IS_0 \propto \frac{M}{\alpha} ln\left(1 + (D^m(N^m, \rho) + D^f(N^f, \rho))^{1-\sigma}\right),$$

where we use different values of N^m and N^f depending on the simulation. For example, $IS_0 = IS(D(N_0^m), D^f(N_0^f))$ corresponds to the status quo, $IS_1 = IS(D(N_0^m), D^f(N_1^f))$ corresponds to eliminating the female influx without equilibrating male entry, and $IS_2 = IS(D(N_1^m), D^f(N_1^f))$ describes the equilibrium male response to the elimination of the female influx. Our main interest is in the proportionate change in IS with the influx, relative to counterfactual environments, or in IS_0 in relation to IS_1 and to IS_2. The welfare effect of a female influx is unambiguously positive absent a male entry response to its elimination. If we allow for a male entry response, however, the female influx need not have a positive effect on innovator surplus. We show in Appendix A.1 that endogenous male entry in this model leads to a $\%\Delta IS$ that is bounded between the no-male-response value and zero, inclusive.

2.3 From model to data

The welfare effects of the female influx depend on both the effects of removing female innovations and the male entry response. In particular, the equilibrium percent change in innovator surplus is

$$\%\Delta IS = 100 \times \left[\frac{\ln(1 + (D_0^f + D_0^m)^{1-\sigma})}{\ln(1 + (D_1^f + D_1^m)^{1-\sigma})} - 1 \right].$$

Each of the components of the welfare effect depends on magnitudes and relationships we can explore directly in the data.

The overall effect has two parts. First, the direct effect of removing the female influx stems from reducing D_0^f to D_1^f. This is large if the female innovations removed have high value. We can explore this by comparing the citations received by female innovations to the citations received by male innovations over time, as the female share rises. For example, how many citations do female innovations receive, per innovation, compared with male innovations? Moreover, as the female share of new innovations has risen, has the share of citations garnered by the new female innovations risen proportionally?

The equilibrium effect also depends, of course, on the male entry response, or the extent to which D^m rises with additional male innovation entry. The counterfactual removal of female innovations raises the expected return (citations) to marginal male innovations. How much equilibrating male entry will fill the gap created by the removal of the female influx depends on three things that determine the shape of the function $D^m(N^m)$: the distribution of male innovation qualities, the substitutability of innovations for one another, and the predictability of innovation quality at the time of entry.

3 Data

Addressing our research question requires data on academic publications and patented inventions, by topic area, the (forward) citations they receive, and the genders of both authors and inventors. We also need to segment publications and patents by subject area. We use already-existing data on patents and their inventors, and we develop a novel approach to ascertaining the genders of academic authors using large language models. We use the academic and patent data to create four analysis datasets.

3.1 Academic dataset construction

Metadata regarding academic articles are obtained from the Microsoft Academic Graph (MAG) (Sinha et al., 2015). Although MAG contains many publication types including preprints, conference proceedings, and books, we restrict our sample to 65.2 million journal articles as these are the most reliably indexed research artifacts (Moed, 2005; Visser et al., 2021). MAG's coverage ends in 2021, so we examine publications between 1965 and 2015 to ensure that every article in our sample has five years to accrue forward citations. Citations are also collected from MAG, limited to those received within five years of the cited article's publication.

It is essential to account for scientific area, but MAG does not contain such an ontology. We leverage the crosswalk from Koffi and Marx (2025), which probabilistically maps journal articles in MAG to OECD fields as well as Web of Science (WoS) subject categories. Most of our analyses, and all of our visualizations, are at the OECD Field of Science and Technology subfield-classification.[8] Finer-grained categories are available from the 251 Web of Science

[8]These are Agricultural and Biological Sciences, Arts and Humanities, Biochemistry, Genetics and Molecular Biology, Business, Management and Accounting, Chemical Engineering, Chemistry, Computer Science, Decision Sciences, Dentistry, Earth and Planetary Science, Economics, Econometrics and Finance, Energy, Engineering, Environmental Science, General, Health Professions, Immunology and Microbiology, Materials Science, Mathematics, Medicine, Multidisciplinary, Neuroscience, Nursing, Pharmacology, Toxicology and Pharmaceutics, Physics and Astronomy, Psychology, Social Sciences, Unknown, and Veterinary.

categories, which we use in robustness tests.

3.1.1 Determining author gender

Our analyses rely on an indicator for whether half or more of an article's authors are female. Let $f_j = F_j/(F_j + M_j + U_j)$ be the female share of authors on article j, where F_j, M_j, and U_j are the numbers of female, male, and unidentified-gender authors of article j. Define the mostly-female indicator $\mathbb{1}_j^f = \mathbb{1}\{f_j \geq 1/2\}$. Because authors of unidentified gender enter the denominator but not the numerator of f_j, the indicator can only under-classify an article as mostly-female, but never over-classify it. Hence, our measures are conservative (lower-bound) estimates of female authorship.

Gender is often inferred from given names ("forenames") using algorithms based on the probability that a given forename belongs to males vs. females (Lerchenmueller and Sorenson, 2018). Even when geography is incorporated into probability tables, however, gender inference algorithms entail substantial errors (Marx, 2024). Our approach uses large language models (LLMs) to produce high-confidence gender inference for authors whose forenames alone yield low-confidence estimates.

The intuition behind this approach is that information incorporated into LLM pre-training might usefully disambiguate ambiguous forenames for certain populations. Consider Chris Forman, a faculty member at Cornell University. Chris' forename is gender-ambiguous and would remain unclassified via algorithm. However, Chris' faculty page at Cornell contains gender-identifying pronouns.[9] This approach is promising for academic authors because many, if not most, have public-facing websites.[10]

[9] Text from his website (https://business.cornell.edu/faculty-research/faculty/cmf257/): Chris Forman is the Peter and Stephanie Nolan Professor of Strategy, Innovation, and Technology at the Charles H. Dyson School of Applied Economics and Management at Cornell University. He received his PhD in managerial economics and strategy from Northwestern University and a BA from the University of Pennsylvania. Prior to joining Cornell, he served on the faculty of the Scheller College of Business, Georgia Institute of Technology, and the Tepper School of Business at Carnegie Mellon University. He has also worked as an associate economist at Lehman Brothers.

[10] Our approach may be less successful for populations that are less publicly visible, such as corporate

We begin with 66,692,537 authors in MAG. For names that can be classified with at least 90% confidence using `gender.api`, we use forename-country-based classification. This corresponds to 22,317,642 distinct authors in our dataset. For the remaining 44,374,895 authors, we proceed as follows. Given that the LLM has a better chance of retrieving the correct gender if the author's affiliation is known, as per the examples above, we look up authors' most recent affiliation (i.e., from their prior publications) in MAG. A current or prior affiliation can be retrieved for 22,742,112 authors. For the remaining 21,632,783 authors lacking current or prior affiliation, we "interpolate" possible affiliations for authors by a) taking the most frequent affiliation of coauthors on a paper where the focal author's affiliation is missing, and b) crosswalking the paper (via DOI) to the Clarivate Web of Science and using the affiliation for the paper. Interpolation covers another 11,734,520 authors, leaving 9,898,263 for whom we lack any affiliation.[11]

To validate this procedure, we manually verified the gender of a random sample of 100 authors classified by either method. We used public information such as institutional webpages, professional profiles, and gender-identifying pronouns or photographs. Manual verification established gender for 66 of the 100 authors. Error rates decline in reported confidence for both the forename-based and the LLM-based classifications. This supports the use of confidence thresholds. Applying our 90-percent confidence rule to the sample assigns gender to 48 of the 100 authors. Of the 42 that could be verified manually, 2 are misclassified. This results in an error rate of 4.8 percent.[12]

patent inventors. Hence, we rely on USPTO data on inventor gender.

[11]We use the 8-billion parameter Llama 3.1 model (4-bit quantized) to retrieve gender using the following prompt: *Determine the gender of these researchers using only internal knowledge: . For each researcher, assign exactly one gender: Male, Female, or Unknown (use Unknown for ambiguous names). Respond strictly with a comma-separated list of exactly genders in the same order as the input names. also include a confidence score from 0-100 for how sure your classification is. Do not include explanations, notes, or any other text. Example: Male/100,Female/50,Unknown/30.* Importantly, the LLM is asked to provide a confidence score for each classification. We treat as known only LLM classifications with a confidence score of 90% or higher.

[12]The hand check also speaks to the direction of misclassification. The two misclassifications run in opposite directions. The LLM-based error assigns male to a female author, and the forename-based error

3.2 Patent data

The patent data for this study consist of 7,626,120 patents granted in the US between 1976 and 2020, along with 106 million citations granted over the same period. For each citing and cited patent, we observe its World Intellectual Property Office (WIPO) technology categories and grant year. We classify patents into 35 technology categories according to their first listed WIPO category.[13] The data are from the USPTO PatentsView repository.

PatentsView contains a field for the gender of each inventor, with about 10% of inventors unclassified due to ambiguous forenames. Given that the vast majority of patents are filed by corporations as opposed to universities, our LLM-based method used for papers is less likely to give us as much of an advantage.

3.3 Analysis datasets

We use the underlying datasets above to create two kinds of datasets for analysis, one of each kind for articles and for patents, yielding four analysis datasets. The first kind is at the micro (article or patent) level. We observe characteristics of each article and patent, including its topical fields and appearance vintage, as well as whether it is mostly-female (half or more of its authors or inventors are female). We also observe the number of citations that each article or patent receives from male or female authors or inventors in each calendar year. We use the micro datasets to estimate the citation demand models. Second, we create aggregate data showing the numbers of mostly-female and other articles or patents by academic or

assigns female to a male author. However, with two errors among 42 verified assignments, the scope for systematic bias in our gender-share measures is in any case limited.

[13]The resulting categories are analysis of biological materials; audio-visual technology; basic communication processes; basic materials chemistry; biotechnology; chemical engineering; civil engineering; computer technology; control digital communication; electrical machinery, apparatus, energy; engines, pumps, turbines; environmental technology; food chemistry; furniture, games; handling; IT methods for management; machine tools; macromolecular chemistry, polymers; materials, metallurgy; measurement; mechanical elements; medical technology; micro-structural and nano-technology; optics; organic fine chemistry; other consumer goods; other special machines; pharmaceuticals; semiconductors; surface technology, coating; telecommunications; textile and paper machines; thermal processes and apparatus; and transport.

technological area and by vintage. These aggregate data also include the numbers of citations received by male and female articles or patents in each vintage and topical field. In some versions of these datasets we also disaggregate to citing year. We use the aggregate datasets for the descriptive analyses.

Tables 1–3 report summary statistics for the four analysis datasets, i.e., the micro and aggregate datasets for each domain. In each table, Panel (a) covers articles and Panel (b) patents.

Table 1 describes the micro data. It consists of a 5 percent random sample of MAG journal articles (N = 3,260,975; vintages 1965–2015) and the census of granted USPTO utility patents (N = 6,892,318; vintages 1976–2020). Female participation is higher in academic science than in patenting: 13.4 percent of articles are mostly-female-authored versus 8.1 percent of patents mostly-female-invented. Citations are heavily right-skewed in both domains. The mean article receives 19.9 citations within five years against a median of 3. The corresponding figures for patents are 3.2 and 1. Because unidentified-gender authors and inventors count toward the denominator, the mostly-female indicator is conservative, particularly for articles.

Table 2 describes the aggregate data. It consists of 2,040 field–vintage pairs for articles (40 OECD fields, 1965–2015) and 1,574 for patents (35 WIPO fields, 1976–2020).[14] The mostly-female share averages 12.1 percent across field–vintage pairs for articles (maximum 38.6 percent) and 7.7 percent for patents (maximum 19.3 percent). The mean share of citations to female-authored articles (14.4 percent) slightly exceeds the female-authored share. For patents the two are comparable (7.2 versus 7.7 percent). Female citers direct 22.2 percent of article citations to female-authored work, roughly twice the male-citer share of 11.9 percent. The patent figures are 12.8 and 6.8 percent.

[14] One field–vintage pair is empty: micro-structural and nano-technology (WIPO field 22) had no granted patents in 1976. Thus the patent panel has 1,574 rather than $35 \times 45 = 1{,}575$ pairs.

Table 3 describes the micro data in the form used for demand estimation. The micro data are reshaped to the cited-work × citing-year level. We obtain 10.5 million article observations (citing years 1980–2015) and 31.8 million patent observations (1990–2020). The dependent variable is the Berry (1994) log inside-share difference, and the log inside share is the endogenous regressor. The excluded instrument is the log number of extant works available to be cited. Each variable is constructed separately using all, female, and male citers, which allows us to estimate demand by citer gender.

4 Descriptive evidence

This section presents descriptive evidence on the female influxes into patenting and academic authorship. First, Section 4.1 documents the influxes overall and by topic area. The remainder of the Section explores the questions implied by the theory outlined in Section 2. Section 4.2 documents the citations received by female innovation over time, relative to male innovation. Section 4.3 presents evidence that additional innovation expands the market by attracting more citations. Section 4.4 explores the predictability of innovation quality at entry. Section 4.5 presents evidence on whether the growth of female innovation crowds out male activity.

4.1 The female influx

For simplicity of exposition, we refer to articles whose authors are mostly female as "female-authored;" and we refer to the complement, the non-majority-female-authored articles, as "male-authored."

Our first questions are simply, how large are the female influxes into academic science and patenting? Panel (a) of Figure 3 shows the share of articles with mostly-female authors. Under 5 percent of the articles published in 1965 have mostly-female authors. This rises

steadily and reaches 20 percent by 2015. Panel (b) of Figure 3 presents the numbers of male and female articles over time, providing a different perspective on the influx. Both female – and male – authorship are growing substantially over the half-century. Female gains appear not, on their face, to come at the expense of male participation. Panels (c) and (d) repeat the exercises for patents. Female-inventor patents rise from under 5 percent in 1976 to over 9 percent in 2020. As with articles, both male and female-inventor patents granted rise over time.

The levels, and increases, vary across academic fields. Panel (a) of Figure 4 shows the mostly-female authorship share in the 1960s and the 2010s by field, ordered by the mostly-female share in the 2010s. The mostly-female share rises in all areas, but the eventual share varies sharply across fields. At the high end, it reaches roughly a third in psychology, sociology, education, and art. By contrast, the mostly-female share remains below 10 percent in physics and various engineering fields. Panel (b) repeats the exercise for patents. Female inventors are most prevalent in biology, biotechnology, and pharmaceuticals, with over 15 percent of inventors. Women are rarer in civil and mechanical engineering, where they account for about six percent of inventors at the end of the sample period.

4.2 The influx and citations to female innovation

4.2.1 Citations per article

How well do female-authored articles get cited, relative to male-authored articles; and how does this change with the influx of female-authored articles? For female-authored, vintage-v articles in area a, the average (eventual, as of the end of the sample period) citations per article are given by $\frac{Q_{avf}}{N_{avf}}$. The analogous ratio for male-authored articles is $\frac{Q_{avn}}{N_{avn}}$. Hence, we can characterize the citation influence of female-authored articles, relative to male-authored articles, by $r_{va} = \left(\frac{Q_{avf}}{N_{avf}} / \frac{Q_{avn}}{N_{avn}}\right)$.

Panel (a) of Figure 5 shows this ratio for academic articles, and Panel (b) repeats the exercise for patents. The two panels tell a similar story. In both, the heavy line describing citations from all innovations hovers close to one: overall, women receive roughly the same number of citations per work as men—a bit more for articles, essentially at parity for patents. And in both, the dashed line lies far above the others, showing that women cite women at much higher rates than men do, with the female-to-male ratio among female citers running between roughly two and four in articles and two to three in patents. The male-citer line stays near one in each panel.

4.2.2 The female innovation and citation shares

If the influx of women into the field is valuable, then female papers and patents should attract citations. We explore this hypothesis via regressions of the form

$$q_{avf} = \phi n_{avf} + \mu_a + \mu_v + \varepsilon_{av}, \tag{4}$$

where q_{avf} is the share of citations attracted by female innovations in topic area a and vintage v. The female share of innovations is n_{avf}, and μ_a and μ_v are area and vintage fixed effects.

Table 4 estimates equation (4) for academic articles (Panels a and c) and patents (Panels b and d). In all panels, citing works are restricted to those whose author/inventor gender is identified.[15] Column (1) pools female and male citers, while columns (2) and (3) use citations from female and male citing works, respectively. Panels (a) and (b) absorb field and vintage fixed effects. Panels (c) and (d) use data by field, vintage, and citing year, augmenting equation (4) with a citing-year fixed effect.

In column (1) of Panels (a) and (b), the coefficient is 1.06 (se = 0.060) for articles and 1.12 (0.124) for patents. As the share of articles or patents by women rises, the share of

[15]Results remain qualitatively unchanged when we also include citing works whose author gender cannot be identified. These results are not reported here but are available upon request.

citations to these works rises proportionally.

Columns (2) and (3) of Panels (a) and (b) repeat the exercise according to the gender of the citing author. Column (2) uses the share of citations from female-authored citing works to female-authored cited works. Column (3) uses citations from male-authored citing works. The female-citing coefficients are larger: 1.06 versus 0.80 for articles, and 2.03 versus 0.96 for patents. (The patent estimate in column (2) is imprecise but tightens to 1.60 with a citing-year fixed effect in Panel d.) Increases in the female-authored share of works raise the share of citations to female authors more for female than for male citing works. This reinforces what was visible in Figure 5, that female authors derive more benefit from female-authored works than male authors do. Panels (c) and (d) deliver nearly identical results. The effects we document in Panels (a) and (b) across vintage remain when we allow for citing behavior that varies across citing year.

The results above suggest that the female influx is valuable. Works by women attract roughly as many citations as those by men, and female authors cite them more than male authors do. That these contributions are valuable does not, by itself, demonstrate that the influx raises the value of research and invention. For instance, contributions by women may displace contributions from men that would have been similarly valuable. We return to this possibility below, both descriptively in Section 4.5 and in the structural model in Section 5.

4.3 Business stealing vs. market expansion

Female contributions come to account for substantial shares of innovations, but this does not guarantee that they add substantially to the value of extant innovations. Additional innovations are valuable if they are differentiated from existing innovations. We could infer that they are valuable, for example, from an increase in total citations resulting from additional innovation. We explore this possible market expansion by asking whether growth in the number of innovations in a given area subsequently raises the number of citations those

works receive. The works providing citations in a given year face a choice set of extant works to cite. The question is whether enlarging that choice set draws in *more* total citations (market expansion) as opposed to merely redistributing a fixed number of citations across more works (i.e., business stealing).

The choice set evolves over time. One simple summary of that choice set is N_{at}, the number of extant innovations in field a appearing up to year t. We ask whether the number of citations from year-t innovations (Q_{at}) directed at extant works in field a (N_{at}) rises faster in fields where the number of works grows more. We examine this via the following regression:

$$Q_{at} = \nu N_{at} + \mu_a + \mu_t + \varepsilon_{at}. \tag{5}$$

We face the complication that we observe articles only since 1965 and patents only since 1976. Hence, $N_{at} = \Sigma_{v=1965}^{t-1} N_{av}$ for articles and $N_{at} = \Sigma_{v=1976}^{t-1} N_{av}$ for patents. Consequently, we experiment with different start dates for the regression.

Table 5 reports estimates of equation (5) in levels (column (1)) and logs (column (2)). Panel (a) covers academic articles. Areas whose extant articles grow attract more citations when accounting for publication year. This market expansion supports the notion that additional articles are valuable. Panel (b) repeats the exercise for patents, with similar results. With log–log elasticities near one for both patents (1.100) and articles (1.119), citations grow roughly one-for-one with the stock rather than being diluted across more works. This suggests low substitutability and therefore that additional innovations would be valuable.

4.4 Innovation quality predictability

The shape of the $D()$ function, and therefore the marginal benefit of entry, depends on the predictability of innovation quality at entry. As discussed in Section 2, the more predictable is quality, the lower the expected citations to marginal entering innovation. We explore this ex ante predictability by running regressions of realized citations on factors known at the time of entry. These include innovators' prior innovations (and the citations they have received), category, and interactions.

Many of the variables enter significantly, in particular the innovator characteristics. Still, the share of variation that is explained – and therefore the predictability of quality at entry – is low. The R-squared measures are 10 and 9 percent for patents and academic articles, respectively. These regressions also provide ex ante quality predictions $\delta_j^{'}$, which we use for the entry order in Section 5.2.

4.5 Displacement

In our theoretical model, the numbers of male and female innovations are connected in equilibrium. The removal of the female influx (reducing N^f and D^f) raises the return to innovation, inducing an increase in N^m. Conversely, the arrival of the female influx would be expected to displace male innovation. We explore this by estimating the following regression model relating the number of articles or patents in a topical area created by men in each vintage to the number of female-created articles or patents appearing in the same area and vintage, controlling for vintage and field fixed effects:

$$N_{av}^m = \psi^f N_{av}^f + \mu_a + \mu_v + \varepsilon_{av}, \tag{6}$$

where N_{av}^m is the number of male works in area a and vintage v, and N_{av}^f is the analogous measure for works created by women. If female-authored works simply displaced male-

authored works one for one, we would expect $\psi^f = -1$.

Table 6 presents estimates of equation (6) for academic articles (Panel a) and patents (Panel b), in levels (Columns (1) and (2)) and in logs (Columns (3) and (4)), with and without field and vintage fixed effects. In both panels, and across all specifications, the estimate of ψ^f is positive and statistically significant. We therefore find no evidence that male contributions are displaced by female-authored articles or female-invented patents. We interpret these results, together with Panels (b) and (d) of Figure 3, as indicating substantial growth in the demand for both male and female scholars and inventors. Moreover, the influx of women has expanded the total amount of research and invention created. In spite of these estimates, we allow for an endogenous male entry response in the structural model below, tempering the magnitude of the equilibrium welfare effect.

5 Welfare estimation

The results above suggest that the female influx into innovation is valuable. Female-authored innovations attract roughly as many citations per article as male-authored articles (at least until recently). Moreover, as the share of innovations by women rises—across vintages within areas—the share of citations garnered by those innovations likewise rises. This increase is roughly proportional, although somewhat lower for citations from male innovations than from female innovations. While the descriptive exercise suggests that the female influx is valuable, it neither gives rise to a quantitative measure of the change in innovator surplus, nor can it provide a consistent way to account for equilibrium counterfactual male entry.

This section quantifies the value of the female influx using simple industrial organization tools. First, we present a model of the demand for innovations to cite. We treat the creators of each year's entering innovations as the "buyers" of existing innovations, which they "buy" via citation. Previously-created innovations at each point in time are the choice set for that

year's innovators. Our goal is to measure the value of the evolving choice set, and we do this both from the standpoint of citing innovators overall, as well as from the respective perspectives of female and male innovators.

The exercise has two components. First, we have a nested logit demand model, which we estimate separately for three groups (overall, male, and female creators) and – in some specifications – separately for each topic area. Second, we explore two counterfactuals, one in which the female influx is removed and another in which male entry is allowed to equilibrate in response to the removed female innovations.

5.1 Demand

We estimate various nested logit models of demand for cited innovations, using different sets of citing innovators. First, we estimate models using all articles or patents, respectively, as potential substitutes for one another. Second, we estimate models separately by innovative area (patent category or academic field). In all cases, we estimate separate variants for citations from patents or articles with female (majority female) vs male (other) inventors or authors. Following Berry (1994), we estimate nested logit demand models via

$$\ln(s_j^t) - \ln(s_0^t) = \gamma^a \mathbb{1}_j^f + \sigma^a \ln\big(s_j^t/(1 - s_0^t)\big) + \mu_v + \xi_j, \tag{7}$$

where s_j^t is the number of citations received by work j (which was published in vintage v and is in field a) during year t, divided by "market size" M^t (which we operationalize as the number of articles or patents appearing in year t times 100). The term $\mathbb{1}_j^f$ is an indicator that is 1 if innovation j is majority-female. μ_v denotes a cited-work vintage fixed effect.

The parameter σ shows the extent to which articles are substitutes. If $\sigma = 1$, then works are close substitutes for one another; if $\sigma = 0$, then this simplifies to the plain logit model; and works are not close substitutes for one another so that additional articles or patents

bring large welfare benefits. It is identified by the extent to which total citations rise across years with the growth in the number of articles or patents to cite. We use the log of the number of articles or patents appearing in each year $\ln(N_t)$ as an instrument for the log of the inside share, $\ln(s_j^t/(1-s_0^t))$. In topic area models we use $\ln(N_t^a)$ as an instrument. We estimate Eq. (7) pooled (a single σ in Tables 7–8) and separately by field a, yielding the substitution parameters σ^a in Figure 6 and the female-innovation preferences γ^a in Figure 7.

Tables 7 and 8 present estimates of the nested logit model in which all innovations are potential substitutes for one another. Table 7 reports first-stage estimates, which are regressions of the log inside shares on simple "BLP instruments," the logs of the numbers of extant innovations. The instruments work well, in the sense that they are highly significant.[16] When more innovations exist, each innovation attracts a smaller share of citations.

Table 8 reports the IV nested logit demand models and, in particular, the substitution parameter σ and the utility differential γ that innovators attach to female innovations. Estimates of σ are reasonable; all are between zero and one. Articles appear to be better substitutes for one another than are patents. Finally, female citers attach positive utility to female innovations, while male citers attach negative utility.

We also estimate field-specific models. That is, we estimate models in which citers can substitute only among innovations in their field. There is some variation across fields, and almost all of the σ estimates fall between zero and one. Figure 6 shows the estimates of the substitution parameters σ^a by topic area, for articles and patents. Figure 7 shows estimates of the γ parameter – the extent to which citers attach additional utility to female innovations – by citer type and topic area. In articles, female citers attach substantially more utility to female innovations than to male innovations. In patents the utility parameters are more similar across citer genders, perhaps because patent examiners impose more discipline on citing behavior.

[16]The relevant F-statistics are squares of the t-statistics in Table 7.

5.2 Counterfactuals

We seek to measure the welfare effects of the female influx. Baseline IS, with the influx, is $IS_0 = IS(D_0^m, D_0^f)$. We are also interested in two counterfactual scenarios. The first simply removes the female influx, reducing the status quo value of female innovations from D_0^f to D_1^f. The ensuing percent change in IS is $\%\Delta IS_1 = \frac{IS_0}{IS(D_0^m, D_1^f)}$. Finally, we are interested in an equilibrium counterfactual in which male entry adjusts to the removal of female innovations, or $\%\Delta IS_2 = \frac{IS_0}{IS(D_1^m, D_1^f)}$.

Removing the influx. We measure the impact of the female influx into academic publishing and patenting by comparing end-of-sample choice sets (2015 for articles and 2020 for patents) to counterfactual choice sets in which female-authored articles or patents are eliminated, at random, to bring articles to their share in each field at the start of the sample (1965-1969 for articles, 1975-1979 for patents). That is, we calculate D_1^f.

Figure 8 reports estimates of the $\%\Delta IS$ based on simple removal of the female influx, using both the pooled and per-field models. The status quo delivers 5.73 (5.24) percent more IS for academic articles and 1.10 (1.37) percent more for patents compared to a simple removal of the female influx using the per-field (pooled) model. Raw effects on female citers are much larger: 14.00 (13.24) percent for articles and 1.71 (1.87) for patents. These "raw" changes in IS differ by the gender of citing innovators. Yet, these estimates rather unrealistically assume no male entry response.

Equilibrium male entry. In general, the marginal expected citations to male entry depends on the $D(N)$ functions. For example, the expected returns to the marginal entering male innovation are given by:

$$E[Q_m(N^m, N^f)] = \frac{D^m(N^m) - D(N^m - 1)}{D^m(N^m) + D^f(N^f)} \frac{(D^m(N^m) + D^f(N^f))^{1-\sigma}}{1 + (D^m(N^m) + D^f(N^f))^{1-\sigma}}. \tag{8}$$

Status quo entry is (N_0^m, N_0^f), so that the cost of male entry, K_0^m can be inferred from the expected return to the N_0^{mth} entering male innovation, or $K_0^m = E[Q_m(N_0^m, N_0^f)]$. Removing the female influx reduces N_0^f to N_1^f, and reduces D_0^f to D_1^f. The elimination of female innovations also raises the expected returns to male innovation to $E[Q_m(N_0^m, N_1^f)] > K_0^m$. The counterfactual level of male entry – N_1^m – equilibrates the model in response to the removal of the female influx, solving $E[Q_m(N_1^m, N_1^f)] = K_0^m$.

Figure 9 illustrates the counterfactual male entry problem. The solid curve shows the expected citations to the marginal entering male innovation at the status quo female entry. The solid vertical line shows status quo male entry, and K_0^m is at the horizontal line. The removal of the female influx raises the expected return to male entry to the dashed curve. Equilibrium male entry – N_1^m – occurs where the dashed curve crosses the K_0^m line.

Given our estimate of D_1^f, solving the entry problem just requires the expected $D^m(N^m)$ function, i.e. a characterization of how the value of male innovations would evolve with additional entry. The D function depends on two sets of the random variables, the realized product qualities (δ_j) as well as their predictions $(\delta_j^{'})$. Hence, the $D(N)$ associated with N^{th} entering product is an integral over $2 \times N$ variables.

We approximate $D(N)$ as follows. First, we use the predicted quality $\delta_j^{'}$ (from the regressions in Section 4.4) for ordering entry. For each of the male innovations $(1, 2, \ldots, N_0^m)$ in the status quo choice set, we calculate $D(N)$. We then approximate based on $D^m(N^m) = A(N^m)^B$, estimated using the top 25 percent of observations before N_0^m. We term this approximation $\hat{D}^m(N^m)$. Using Equation (8) we can then calculate K_0^m, N_1^m, and the associated $\%\Delta IS$.

Figure 8 shows the aggregate results with equilibrating male entry. The top panel shows results for academic articles. The female influx raises overall IS by 2.29 (2.25) percent using the per-field (pooled) model. Male authors benefit by 2.02 (2.10) percent while female authors derive a benefit of 3.39 (2.84) percent. The bottom panel shows results for patent

inventors. Inventors overall derive a 0.82 percent increase in surplus, with nearly identical benefits across inventor gender using the per-field model. Using the pooled model, the overall benefit to patent inventors is 0.49 percent; and the pooled model somewhat oddly delivers a smaller benefit to female patent inventors.

Figure 10 shows results by innovative field from the per-field models. Panel (a) shows the $\%\Delta IS$ for academic fields, with all authors represented by black bars. The overall benefit is highest for educational science, law, other humanities, and health sciences. Overall benefits are smallest for physics and astronomy and various kinds of engineering. Benefits to female authors are systematically larger than benefits to male authors. Panel (b) shows patent results. The largest overall benefits accrue in pharmaceuticals, basic materials, and various chemistry areas. The smallest benefits arise in optics, medical technology, and digital communications. The differences between the impacts on male and female citers are larger for articles than for patents, echoing the similarity of the male and female citer utility parameters in Panel (b) of Figure 7.

5.3 Magnitudes and marginal benefits

The $\%\Delta IS$ measures quantify the benefits of the total influxes into fields. Because of the larger influx into articles than patents, $\%\Delta IS$ is of course larger for articles than for patents. It is also of interest to understand the marginal benefit of female innovation activity by field. To this end we calculate $\frac{\Delta IS}{\Delta N^f}$ by field. Figure 11 reports these estimates, with articles in Panel (a). Some of the fields with low total benefits have high marginal benefits, suggesting large potential impacts of extending women's involvement. For example, chemical sciences, medical engineering, and biological sciences have low total benefits (see Figure 10) but large marginal benefits in Figure 11. Panel (b) repeats the exercise for patents. Here, too, some of the fields with high marginal benefits have low total benefits. Examples include chemical engineering, materials and metallurgy, and thermal processes and apparatus. The estimates

of high marginal benefits in fields with still-lagging female participation suggest the benefit of a continued female influx.

6 Conclusion

The past half century has seen a substantial female influx into patenting and, especially, into academic research. The female-authored share of new academic articles quadrupled from 5 to 20 percent in the last half century; and the female share of patent inventors rose from 5 to nearly 10 percent over a similar period. We offer descriptive findings suggestive that the female influx is valuable. The growth in female-authored work is accompanied by roughly proportionate growth in the share of eventual citations garnered by female-authored work, and female-authored work attracts roughly as many citations, per patent or article, as male-authored work. Women cite patents and articles by fellow women at elevated rates. Moreover, additional articles and patents appear to "expand the market" by attracting additional citations and not just diverting citations from other works. We see no evidence that the female influx displaces research activity by men. We use an explicit equilibrium welfare framework to measure the proportionate impact of the female influx on the welfare of researchers. Even allowing for an equilibrium male response to the elimination of the female influx, the female influx brings substantial benefits to researchers overall, with larger benefits to female innovators. Finally, the marginal benefits of additional female involvement appear large in a variety of areas, including STEM areas, suggesting benefits from continuing the trend toward gender-inclusive innovation.

References

Aguiar, L. and J. Waldfogel (2018): "Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music," *Journal of Political Economy*, 126, 492–524.

Altonji, J., P. Arcidiacono, and A. Maurel (2016): "Chapter 7 - The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects," Elsevier, vol. 5 of *Handbook of the Economics of Education*, 305–396.

Bell, A., R. Chetty, X. Jaravel, N. Petkova, and J. Van Reenen (2019): "Who becomes an inventor in America? The importance of exposure to innovation," *The Quarterly Journal of Economics*, 134, 647–713.

Berry, S. T. (1994): "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, 242–262.

Blank, D. M. and G. J. Stigler (1957): "Appendix B: Census Data on Number of Engineers and Chemists, 1890–1950," in *The Demand and Supply of Scientific Personnel*, Princeton, NJ: National Bureau of Economic Research, 143–155.

Bleske-Rechek, A. and K. Browne (2014): "Trends in GRE scores and graduate enrollments by gender and ethnicity," *Intelligence*, 46, 25–34.

Ceci, S. J. and W. M. Williams (2011): "Understanding current causes of women's underrepresentation in science," *Proceedings of the National Academy of Sciences*, 108, 3157–3162.

Costa, D. L. (2000): "From Mill Town to Board Room: The Rise of Women's Paid Labor," *The Journal of Economic Perspectives*, 14, 101–122.

Ding, W. W., F. Murray, and T. E. Stuart (2006): "Gender differences in patenting in the academic life sciences," *Science*, 313, 665–667.

Ducor, P. (2000): "Coauthorship and coinventorship," *Science*, 289, 873–875.

Einiö, E., J. Feng, and X. Jaravel (2019): "Social push and the direction of innovation," *Available at SSRN 3383703*.

Fechner, H. and M. S. Shapanka (2018): "Closing diversity gaps in innovation: Gender, race, and income disparities in patenting and commercialization of inventions," *Technology & Innovation*, 19, 727–734.

Goldin, C. (2014): "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 104, 1091–1119.

Harhoff, D., F. Narin, F. M. Scherer, and K. Vopel (1999): "Citation Frequency and the Value of Patented Inventions," *The Review of Economics and Statistics*, 81, 511–515.

Huang, J., A. J. Gates, R. Sinatra, and A.-L. Barabási (2020): "Historical comparison of gender inequality in scientific careers across countries and disciplines," *Proceedings of the National Academy of Sciences*, 117, 4609–4616.

Hyde, J. S., S. M. Lindberg, M. C. Linn, A. B. Ellis, and C. C. Williams (2008): "Gender Similarities Characterize Math Performance," *Science*, 321, 494–495.

Hyde, J. S. and M. C. Linn (2006): "Gender Similarities in Mathematics and Science," *Science*, 314, 599–600.

Iaria, A., C. Schwarz, and F. Waldinger (2022): "Gender Gaps in Academia: Global Evidence over the Twentieth Century," Discussion Paper 17422, Centre for Economic Policy Research.

Jiang, X. (2021): "Women in STEM: Ability, preference, and value," *Labour Economics*, 70, 101991.

Kim, S. D. and P. Moser (2021): "Women in science. Lessons from the baby boom," Tech. rep., National Bureau of Economic Research.

Koffi, M. and M. Marx (2025): "Cassatts in the Attic: Is There a Gender Gap in the Commercialization of Science?" *American Economic Journal: Applied Economics*, forthcoming.

Koning, R., S. Samila, and J.-P. Ferguson (2021): "Who do we invent for? Patents by women focus more on women's health, but few women get to invent," *Science*, 372.

Larivière, V., C. Ni, Y. Gingras, B. Cronin, and C. Sugimoto (2013): "Bibliometrics: Global gender disparities in science," *Nature*, 504, 211–213.

Lerchenmueller, M. J. and O. Sorenson (2018): "The gender gap in early career transitions in the life sciences," *Research Policy*, 47, 1007–1017.

Mammen, K. and C. Paxson (2000): "Women's Work and Economic Development," *Journal of Economic Perspectives*, 14, 141–164.

Marx, M. (2024): "Can we win the "names game"? Tactics for large-scale demographic research in entrepreneurship and innovation," *Industry and Innovation*, 1–9.

Moed, H. F. (2005): *Citation Analysis in Research Evaluation*, Dordrecht: Springer.

Murray, F. and L. Graham (2007): "Buying science and selling science: Gender differences in the market for commercial science," *Industrial and Corporate Change*, 16, 657–689.

National Academy of Sciences (2006): *Beyond bias and barriers: Fulfilling the potential of women in academic science and engineering*, National Academies Press.

Pidgeon, M. E. (1954): "Changes in Women's Occupations, 1940–1950," Women's Bureau Bulletin 253, U.S. Department of Labor, Women's Bureau, Washington, DC.

Shen, H. (2013): "Mind the gender gap," *Nature*, 495, 22–24.

Sinha, A., Z. Shen, Y. Song, H. Ma, D. Eide, B.-J. Hsu, and K. Wang (2015): "An overview of Microsoft Academic Service (MAS) and applications," in *Proceedings of the 24th international conference on world wide web*, 243–246.

Tartari, V. and A. Salter (2015): "The engagement gap: Exploring gender differences in University–Industry collaboration activities," *Research Policy*, 44, 1176–1191.

Thursby, J. G. and M. C. Thursby (2005): "Gender patterns of research and licensing activity of science and engineering faculty," *The Journal of Technology Transfer*, 30, 343–353.

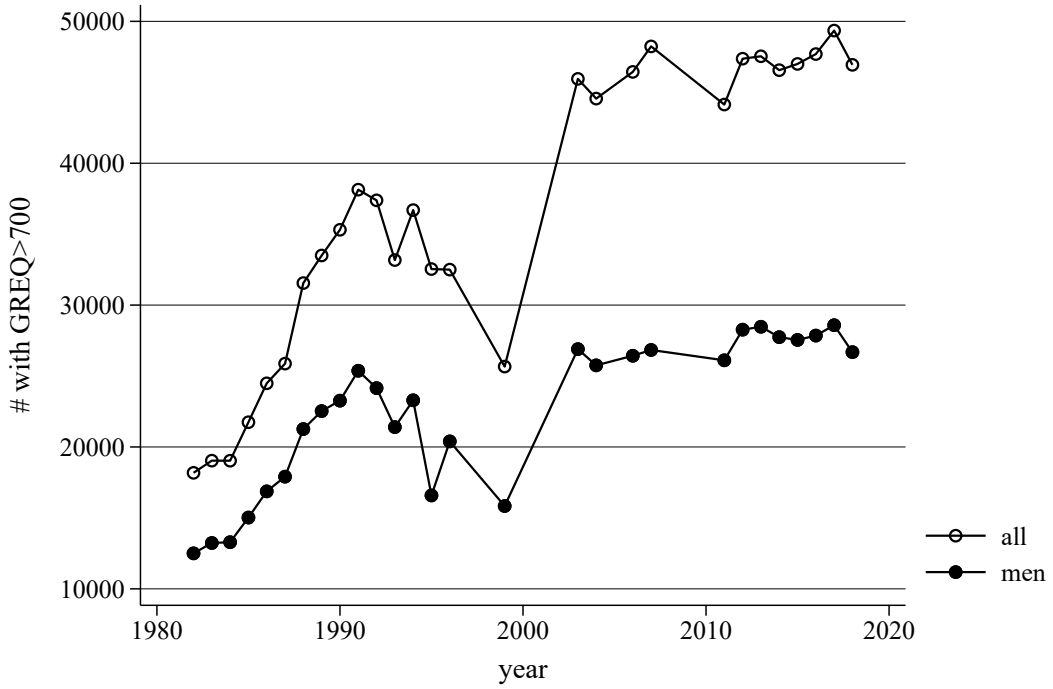
Visser, M., N. J. van Eck, and L. Waltman (2021): "Large-scale comparison of bibliographic data sources: Scopus, Web of Science, Dimensions, Crossref, and Microsoft Academic," *Quantitative Science Studies*, 2, 20–41.

Waldfogel, J. (2025): "The Welfare Effects of Gender-Inclusive Intellectual Property Creation: Evidence from Books," *Journal of Political Economy*, 133, 2229–2264.

Whittington, K. B. and L. Smith-Doerr (2005): "Gender and commercial science: Women's patenting in the life sciences," *The Journal of Technology Transfer*, 30, 355–370.

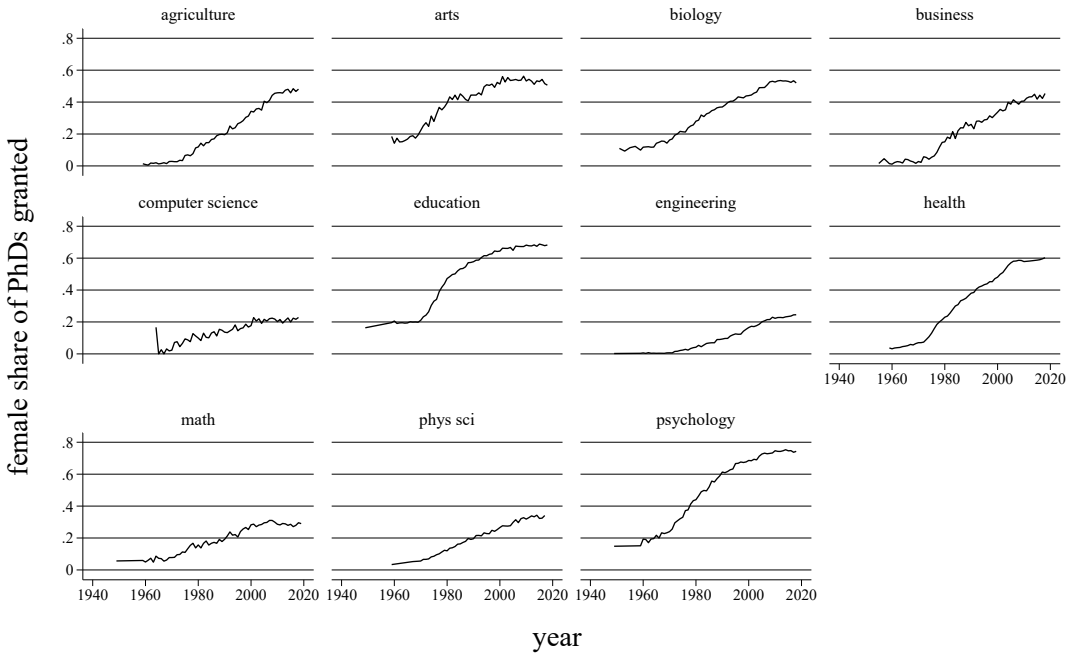
7 Tables and figures

Figure 1: Male and total persons with GRE-Q > 700



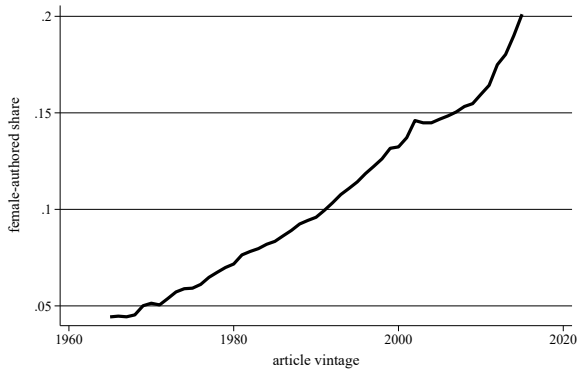
Notes: The number of quantitative GRE test takers with scores above 700 using the old scoring standard. Computed from reported means and standard deviations of scores, assuming normality.

Figure 2: Female share of PhDs earned

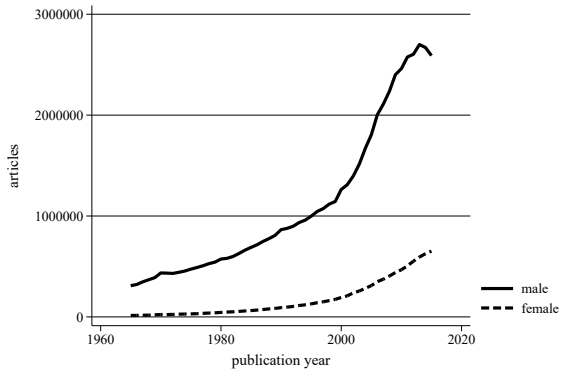


Notes: Each panel plots the female share of U.S. doctor's degrees conferred in the indicated field, by year, 1949–2020. Source: U.S. Department of Education, National Center for Education Statistics, Digest of Education Statistics, Tables 325.10–325.95 (2019–2021 editions).

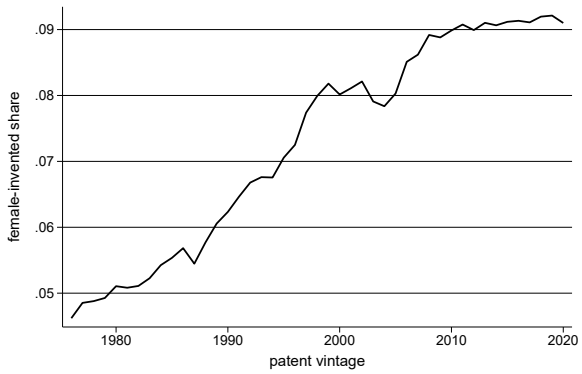
Figure 3: Male and female articles and patents



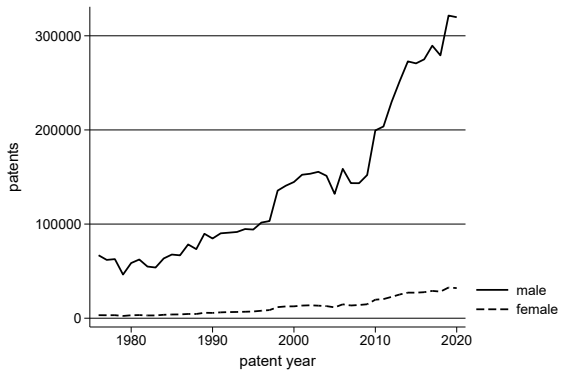
(a) female authorship share



(b) male and female articles



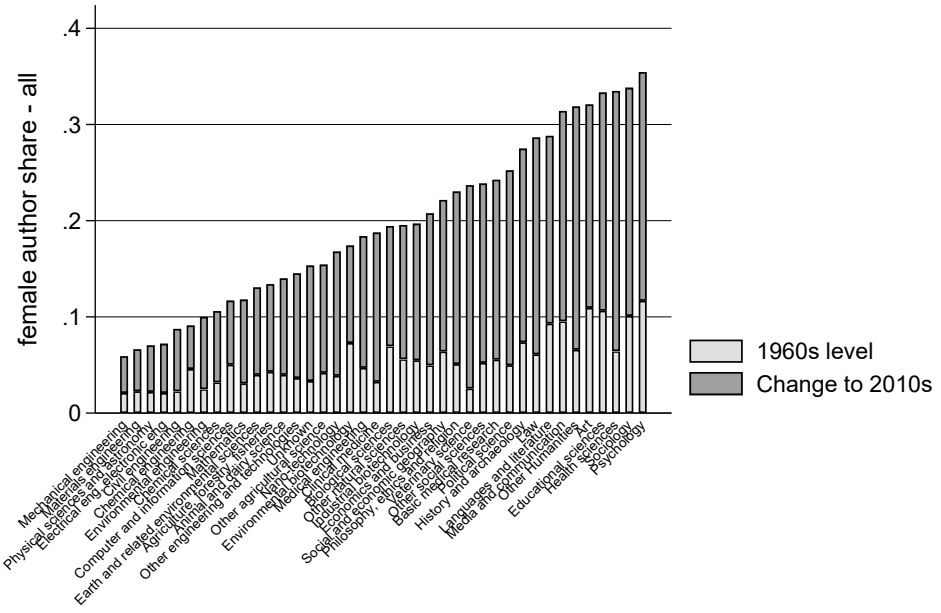
(c) female patent inventor share



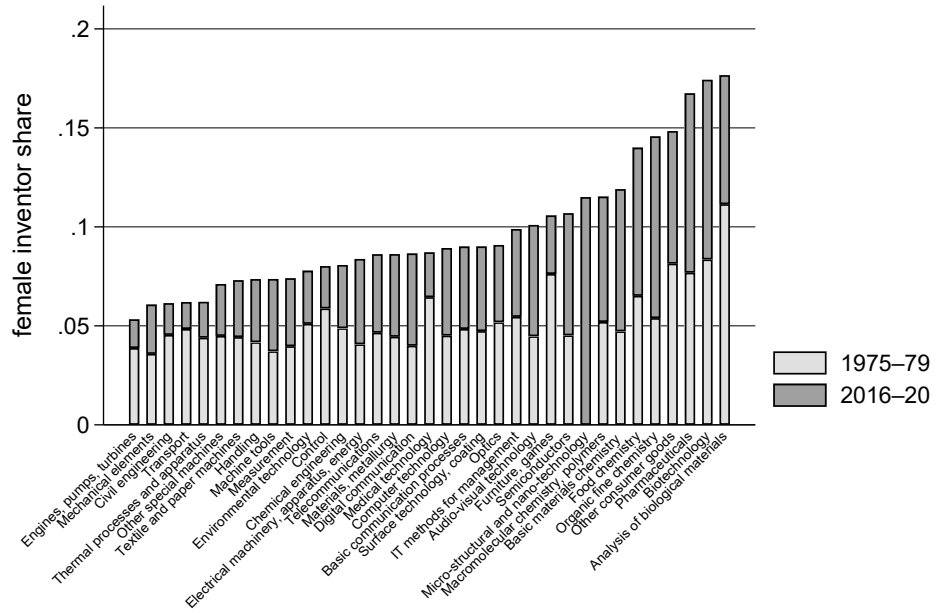
(d) male and female patents

Notes: Panel (a) shows the share of new articles by publication year with mostly-female authors. Panel (b) shows the absolute number of new articles by year from majority-female author teams and the remainder (male). Panel (c) shows the share of new patents by year with mostly-female inventors. Panel (d) shows the absolute number of new patents by year from majority-female inventor and the remainder (male).

Figure 4: Female author and inventor shares by topic, 1960/70s vs 2010s



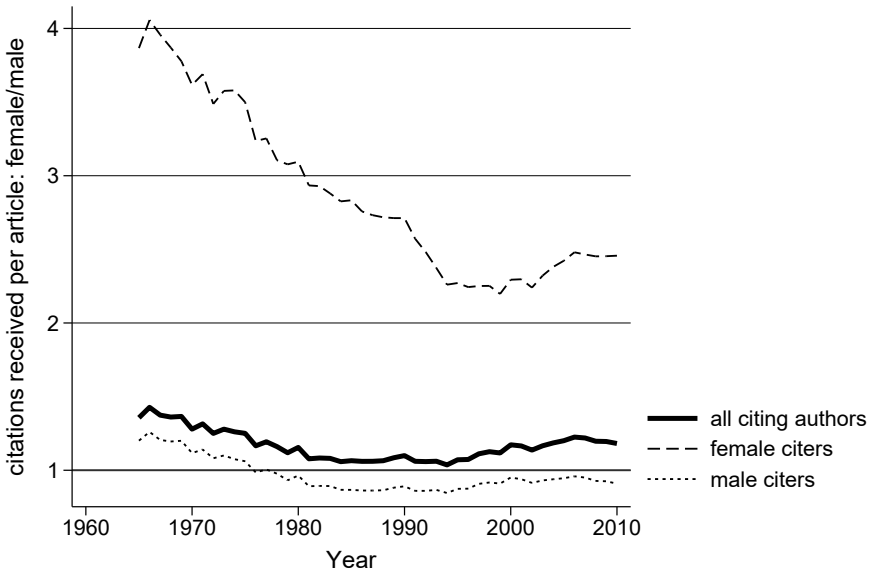
(a) articles



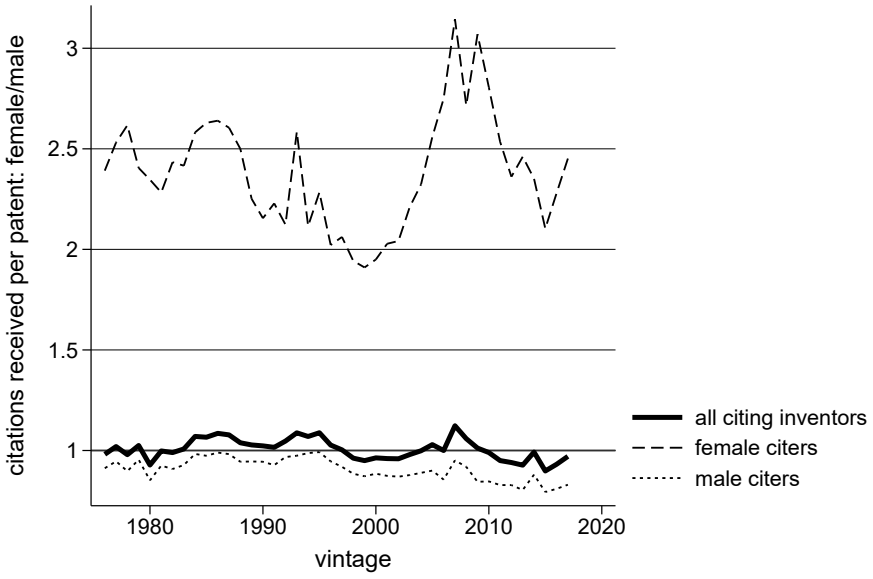
(b) patents

Notes: Shares of articles (a) and patents (b) in each field that are mostly-female-authored in the 1960/70s (light portion of bar, 1965-1969 for articles, 75-79 for patents) vs 2010s (full bar, 2010-2019).

Figure 5: Relative citations per article and patent: F/M



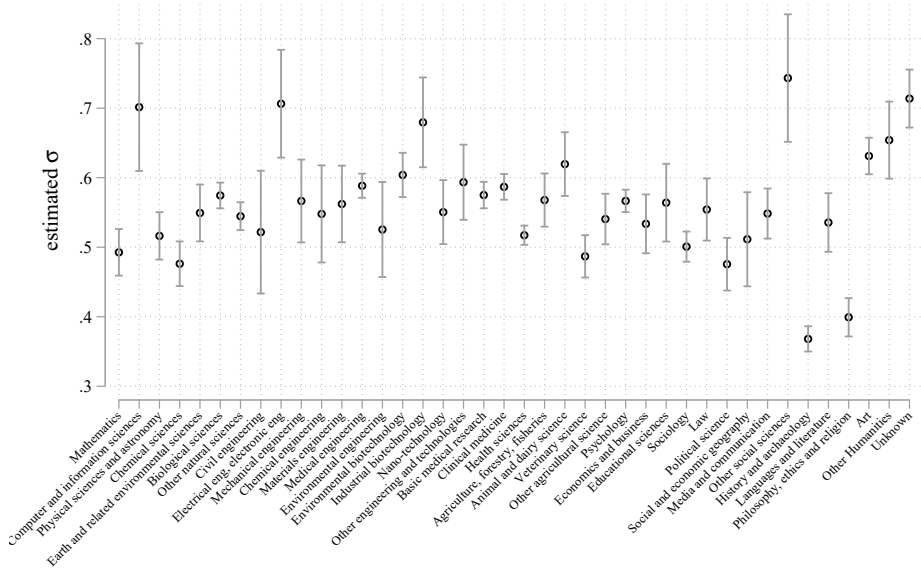
(a) articles



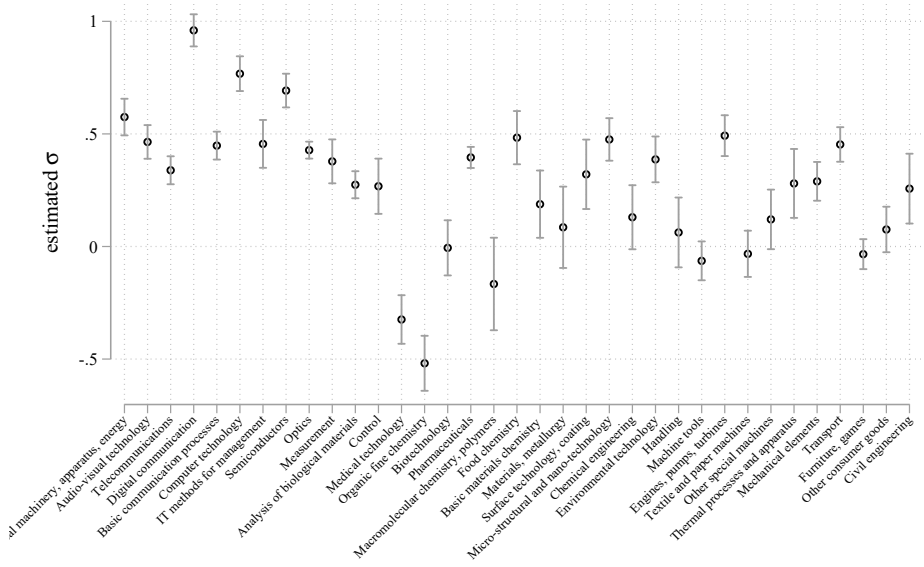
(b) patents

Notes: The ratio of citations per female-authored articles to the citations per non-female-authored articles (a) or patent (b).

Figure 6: Substitution parameter σ by field



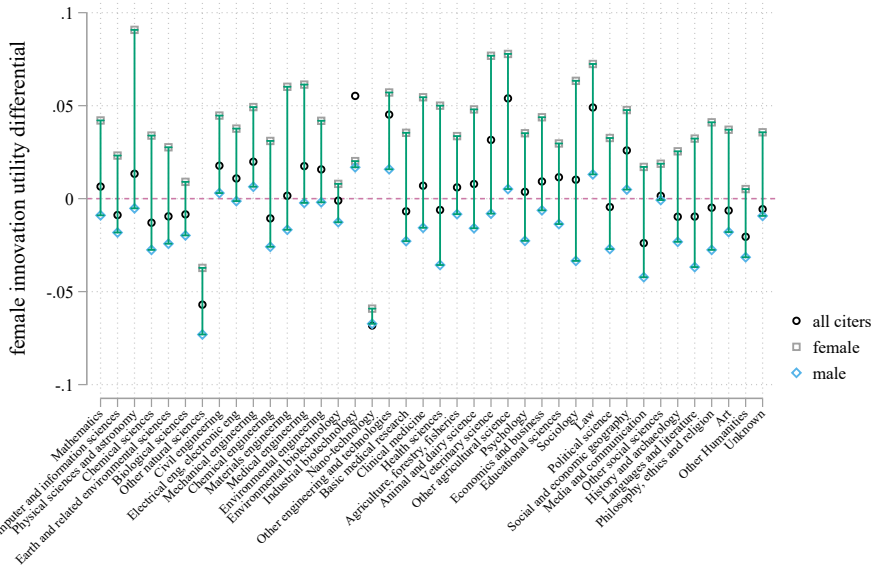
(a) articles



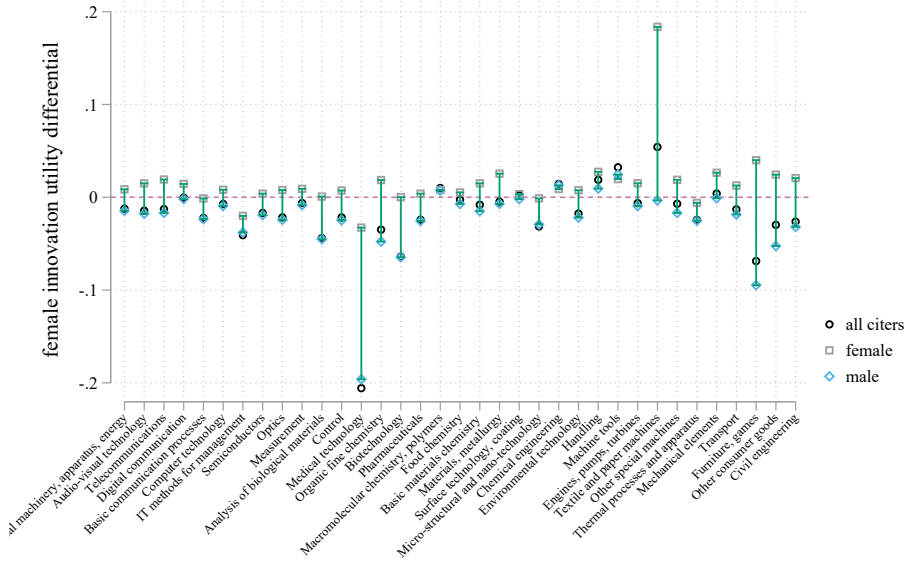
(b) patents

Notes: Markers are the estimated field-specific nested-logit substitution parameter σ^a from Eq. (7), estimated separately by field, with confidence intervals (whiskers). Panel (a) articles, panel (b) patents. A larger σ^a indicates closer substitutability among works within the field.

Figure 7: Preference for female innovations by field



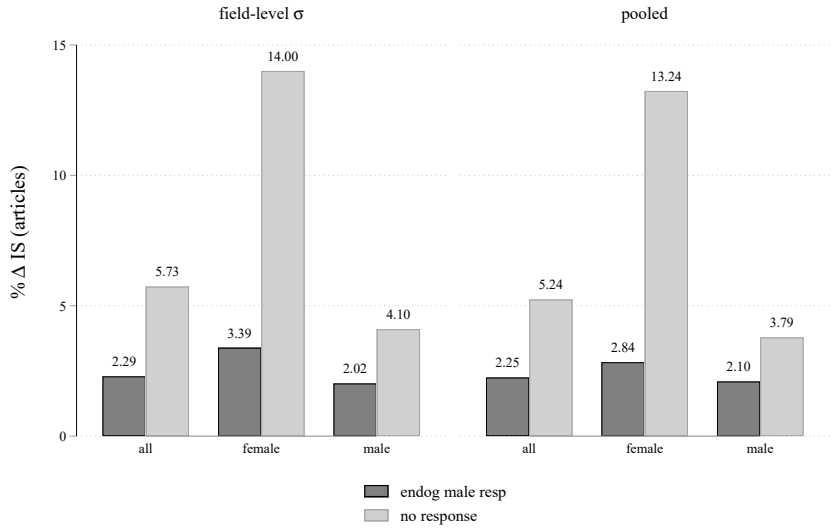
(a) articles



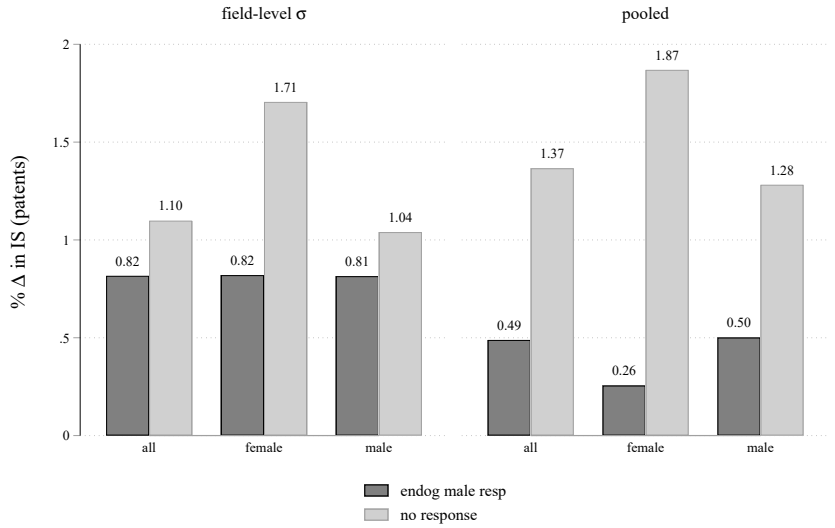
(b) patents

Notes: Markers are the estimated field-specific preference for female innovations γ^a from Eq. (7), estimated separately by field and by citer group (all, female, and male citers), with confidence intervals (whiskers). Panel (a) articles, panel (b) patents. Values above (below) the dashed zero line indicate that citers attach positive (negative) additional utility to female-authored (-invented) works.

Figure 8: Aggregate counterfactual results



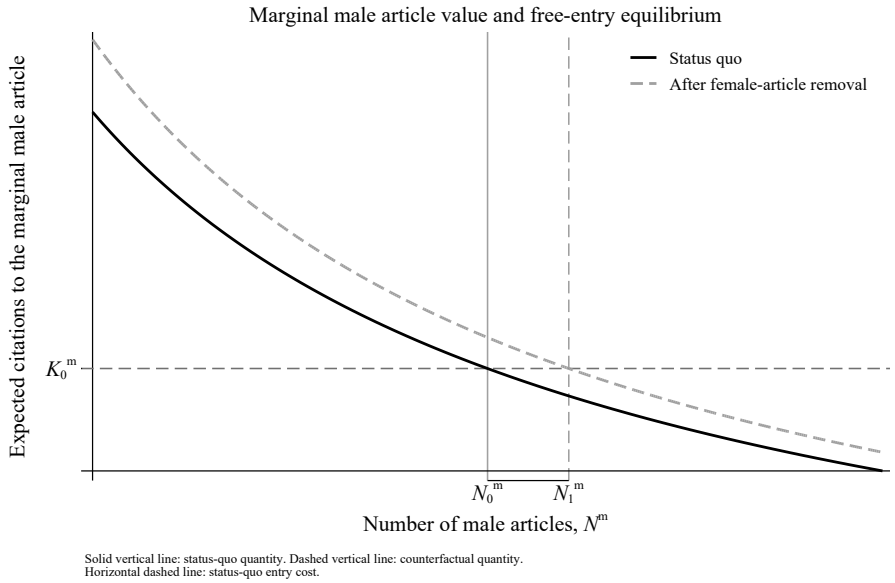
(a) articles



(b) patents

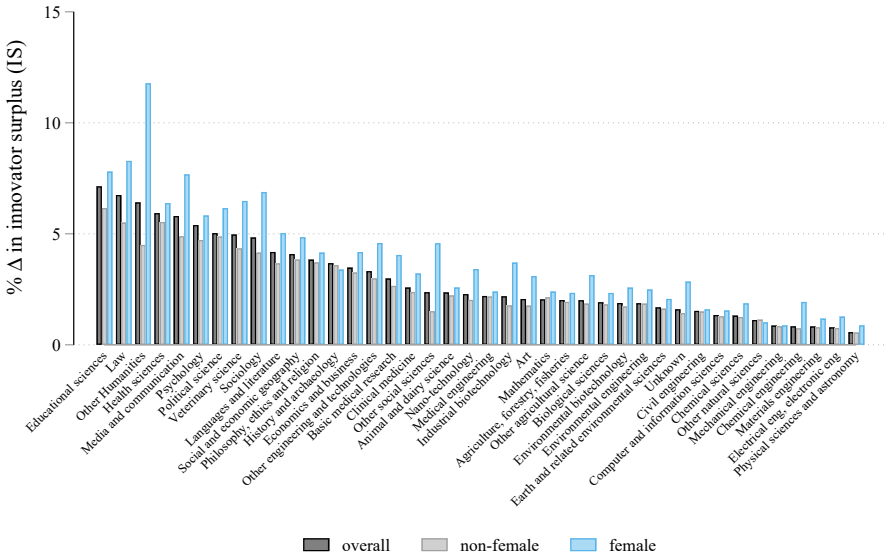
Notes: Bars report the percent change in innovator surplus (%ΔIS) from the female influx, by creator group (all, female, male), under two counterfactuals: holding male entry fixed ("no response") and allowing an endogenous male entry response ("endog male resp"). Within each bar, the lighter shaded area shows the "no response" counterfactual and the darker shaded area the "endog male resp" counterfactual. Results are shown for the field-level σ and pooled demand specifications. Panel (a) articles, panel (b) patents.

Figure 9: Status quo and counterfactual equilibrium entry of male-authored articles

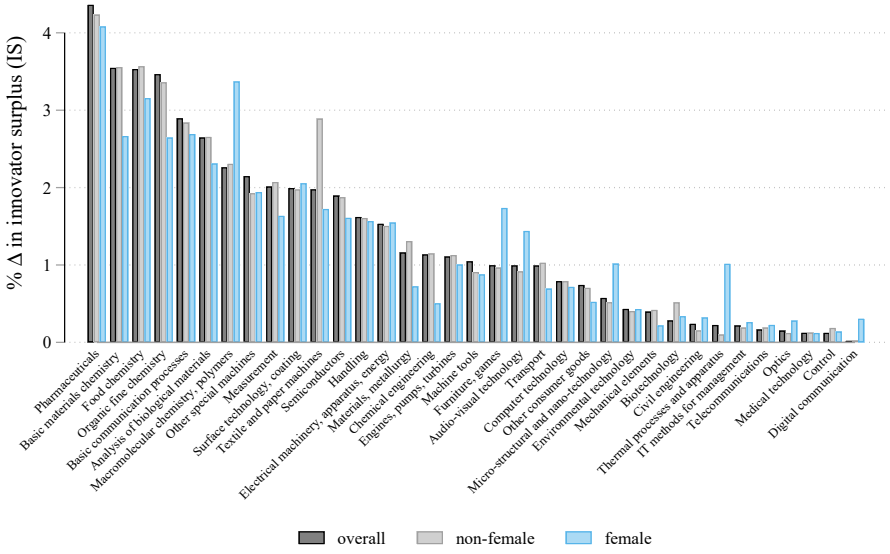


Notes: Schedules show the expected citations to the marginal entering male article. The solid line is based on status quo female entry. The dashed line reflects removal of the female influx.

Figure 10: Field-specific counterfactual results



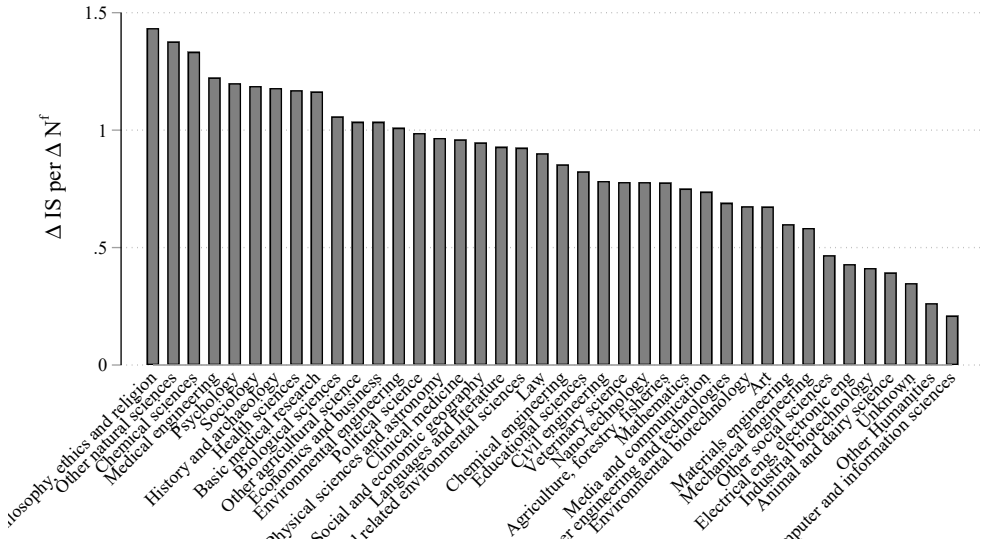
(a) articles



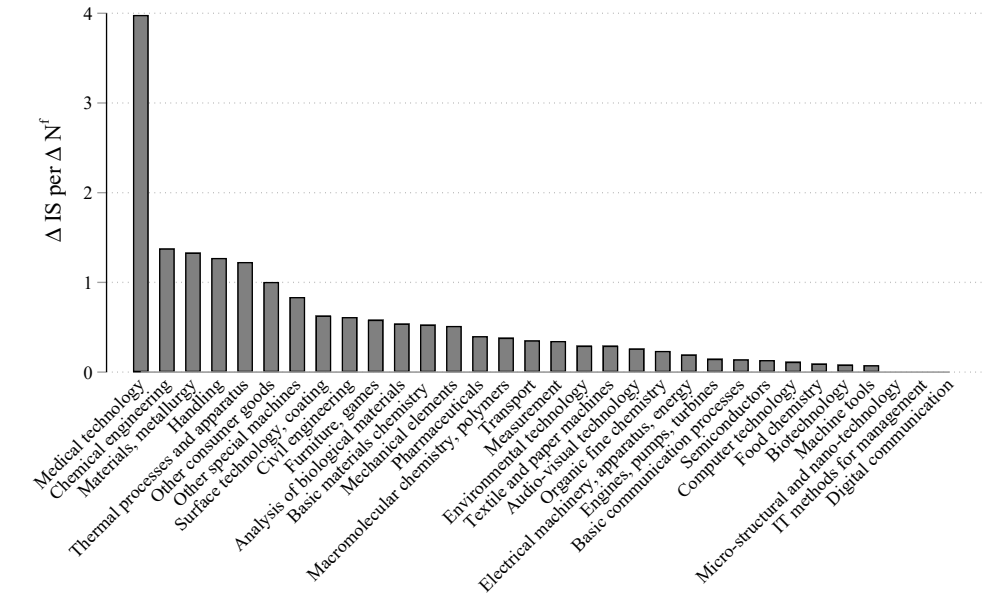
(b) patents

Notes: Bars report the equilibrium percent change in innovator surplus (%ΔIS) from the female influx, by field and creator group (overall, non-female, female), allowing an endogenous male entry response, from the field-level demand specification. Fields are ordered by the overall effect. Panel (a) articles, panel (b) patents.

Figure 11: Marginal benefit of female entry by field



(a) articles



(b) patents

Notes: Bars report the marginal benefit of female entry, i.e., the increase in innovator surplus per additional mostly-female work by field, ordered from largest to smallest. Panel (a) articles, panel (b) patents.

Table 1: Summary statistics, innovation (micro) level: articles and patents

	Mean	SD	Min	Median	p90	Max
Panel A: Academic articles (5% sample, N = 3,260,975)						
Mostly-female-authored	0.134	0.341	0.000	0.000	1.000	1.000
Share of female authors	0.138	0.272	0.000	0.000	0.500	1.000
Share of male authors	0.494	0.405	0.000	0.500	1.000	1.000
Share of unidentified-gender authors	0.367	0.414	0.000	0.200	1.000	1.000
Publication year (vintage)	1999.8	13.1	1965	2004	2014	2015
Citations received within 5 years of publication	19.852	106.954	0	3	47	108781
Panel B: Patents (N = 6,892,318)						
Mostly-female-invented	0.081	0.273	0.000	0.000	0.000	1.000
Share of female inventors	0.086	0.212	0.000	0.000	0.333	1.000
Share of male inventors	0.804	0.308	0.000	1.000	1.000	1.000
Share of unidentified-gender inventors	0.109	0.240	0.000	0.000	0.500	1.000
Grant year (vintage)	2005.1	12.0	1976	2008	2019	2020
Citations received within 5 years of grant	3.201	9.860	0	1	8	1432

Notes: Panel A reports a 5% random sample of MAG journal articles, vintages 1965–2015 (full corpus, cited and uncited). Panel B reports the full census of USPTO PatentsView granted utility patents, vintages 1976–2020 (cited and uncited; withdrawn patents excluded); vintage is the grant year. The citation row counts forward citations received within five years of the vintage year, so a patent uncited within that window—including one never cited—appears with zero five-year citations. Mostly-female-authored (-invented) indicates that half or more of an article's authors (a patent's inventors) are female; unidentified-gender authors and inventors count toward the denominator.

Table 2: Summary statistics, field–vintage level (estimation sample): articles and patents

	Mean	SD	Min	Median	p90	Max
Panel A: Academic articles (N = 2,040 field–vintage cells)						
Mostly-female-authored share of articles	0.121	0.080	0.000	0.099	0.240	0.386
Share of citations to female-authored	0.144	0.104	0.000	0.115	0.303	0.746
Share: female-to-female citations	0.222	0.129	0.000	0.197	0.415	0.838
Share: male-to-female citations	0.119	0.079	0.000	0.099	0.235	0.606
Vintage	1990.0	14.7	1965	1990	2010	2015
Female-authored articles (000s)	4.285	10.662	0.000	1.044	10.485	151.094
Male-authored articles (000s)	27.685	53.646	0.027	9.392	73.908	575.376
Citations to female-authored (000s)	95.400	252.733	0.000	17.119	205.453	2153.033
Citations to male-authored (000s)	539.842	1194.006	0.282	124.410	1257.907	9253.612
Panel B: Patents (N = 1,574 field–vintage cells)						
Mostly-female-invented share of patents	0.077	0.032	0.000	0.070	0.120	0.193
Share of citations to female-invented	0.072	0.033	0.000	0.065	0.116	0.468
Share: female-to-female citations	0.128	0.094	0.000	0.101	0.241	0.739
Share: male-to-female citations	0.068	0.029	0.000	0.062	0.106	0.480
Vintage	1998.0	13.0	1976	1998	2016	2020
Female-invented patents (000s)	0.354	0.476	0.000	0.203	0.766	4.449
Male-invented patents (000s)	4.025	4.733	0.001	2.812	8.464	45.422
Citations to female-invented (000s)	4.187	5.904	0.000	2.321	9.349	53.322
Citations to male-invented (000s)	57.565	85.877	0.001	33.465	120.471	861.643

Notes: Articles are aggregated to academic field (40 OECD fields) by vintage; patents to WIPO technology field (35 fields) by vintage. One field–vintage pair is empty: micro-structural and nano-technology (WIPO field 22) had no granted patents in 1976. So the patent panel has 1,574 rather than 1,575 cells. Counts cover all articles published (Panel A) / utility patents granted (Panel B) in each field and vintage, and are reported in thousands (000s), as are citations. The female-to-female (male-to-female) citation-share row is the share of citations from female-authored/invented (male-authored/invented) citing works that go to female-authored/invented works; these rows are computed over cells with positive citations from the relevant citer group.

Table 3: Summary statistics, demand estimation sample: articles and patents

	Mean	SD	Min	Median	p90	Max
Panel A: Academic articles (N = 10,517,480)						
Female-authored article	0.134	0.341	0.000	0.000	1.000	1.000
Citing year	2004.6	8.9	1980	2007	2014	2015
Vintage	1994.9	11.6	1965	1996	2009	2015
Log articles published in citing year	14.530	0.490	13.334	14.725	15.008	15.008
Dependent variable (all citers)	-18.458	0.887	-19.605	-18.536	-17.261	-11.173
Dependent variable (female citers)	-18.890	0.675	-19.612	-18.918	-17.986	-12.818
Dependent variable (male citers)	-18.517	0.861	-19.607	-18.587	-17.344	-11.304
Log inside share (all citers)	-13.275	1.072	-14.868	-13.392	-11.840	-6.326
Log inside share (female citers)	-12.114	1.005	-13.340	-12.330	-10.636	-6.261
Log inside share (male citers)	-13.121	1.030	-14.624	-13.228	-11.751	-6.211
Panel B: Patents (N = 31,814,328)						
Female-inventor patent	0.071	0.257	0.000	0.000	0.000	1.000
Citing year	2009.3	7.9	1990	2011	2019	2020
Vintage	1997.4	9.9	1976	1998	2011	2020
Log patents granted in citing year	12.277	0.388	11.412	12.322	12.771	12.778
Dependent variable (all citers)	-16.123	0.811	-17.167	-16.336	-15.012	-10.626
Dependent variable (female citers)	-16.700	0.563	-17.368	-16.659	-15.966	-12.108
Dependent variable (male citers)	-16.154	0.802	-17.185	-16.363	-15.035	-10.668
Log inside share (all citers)	-14.397	0.979	-15.748	-14.519	-13.069	-8.828
Log inside share (female citers)	-12.250	0.786	-13.245	-12.232	-11.112	-7.828
Log inside share (male citers)	-14.337	0.969	-15.668	-14.434	-12.995	-8.779

Notes: The unit of observation is a cited article/patent × citing-year cell. The dependent variable is the Berry (1994) log inside-share difference $\ln(s_{jt}) - \ln(s_{0t})$; the log inside share, $\ln(s_{jt}/s_t)$, is the nested-logit regressor. The log number of articles published (patents granted) in the citing year is the excluded instrument; it scales the market size $M_t = 100\,N_t$, where N_t is the number of works entering in year t. The article panel covers citing years 1980–2015; the patent panel covers citing years 1990–2020. Statistics are computed over nonmissing observations: the female- (male-) citer variables are defined only for cells with positive citations from female (male) citers.

Table 4: Regressions of q_{avf} on n_{avf}: articles and patents

	(1) Both citers	(2) Female citers	(3) Male citers
Panel A. Articles, field × vintage			
mostly-female share	1.062***	1.055***	0.803***
	(0.0598)	(0.0722)	(0.0506)
Observations	2040	2040	2040
R-squared	0.885	0.860	0.858
Panel B. Patents, field × vintage			
mostly-female share	1.118***	2.029***	0.961***
	(0.124)	(0.507)	(0.155)
Observations	1574	1571	1574
R-squared	0.768	0.750	0.718
Panel C. Articles, field × vintage × citing year			
mostly-female share	1.101***	1.189***	0.838***
	(0.0614)	(0.0858)	(0.0572)
Observations	58048	56348	57997
R-squared	0.621	0.511	0.528
Panel D. Patents, field × vintage × citing year			
mostly-female share	1.117***	1.596***	0.997***
	(0.175)	(0.119)	(0.191)
Observations	40811	38478	40690
R-squared	0.162	0.313	0.132

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Constant not reported.

Notes: Dependent variable is the share of citations garnered by mostly-female-authored articles (Panels A, C) / mostly-female-invented patents (Panels B, D), regressed on the mostly-female share n_{avf}. In all panels the citing works are restricted to those whose author/inventor gender is identified: column (1) pools female and male citers, while columns (2) and (3) use citations *from* female- and male-authored/invented citing works, respectively. Panels A–B absorb field and vintage fixed effects; Panels C–D add a citing-year fixed effect. Standard errors are clustered by field in Panels A–B, two-way by field and vintage in Panel C, and by vintage in Panel D: the two-way cluster-robust variance matrix is non-positive-semi-definite for the male-citers specification of Panels A–B and the pooled-citer specification of Panel D, and wild-cluster-bootstrap inference under two-way clustering is nearly identical where the comparison is available.

Table 5: Market expansion regressions: articles and patents

	(1) levels	(2) logs
Panel A. Articles		
extant articles in t	1.056***	
	(0.00975)	
log extant articles in t		1.119***
		(0.0319)
Observations	2000	2000
R-squared	0.932	0.995
Panel B. Patents		
extant patents in t	1.262***	
	(0.0504)	
log extant patents in t		1.101***
		(0.0142)
Observations	1574	1574
R-squared	0.642	0.986

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Constant not reported. All columns include field and citing-year fixed effects.

Notes: Dependent variable is the number of citations made in year t to articles (Panel A) / patents (Panel B) of the respective field. The explanatory variable of interest is the number of articles (Panel A) / patents (Panel B) available to be cited in year t. These include all articles published from 1965 (Panel A) / patents granted from 1976 (Panel B) until year $t-1$, by field. Column (1) uses levels, column (2) logs of both variables. Conventional standard errors are reported in all columns.

Table 6: Displacement regressions: articles and patents

	(1) levels	(2) levels	(3) logs	(4) logs
Panel A. Articles				
female-authored articles	4.491***	3.368***		
	(0.0503)	(0.0366)		
log female-authored articles			0.812***	0.549***
			(0.00787)	(0.0149)
Field & vintage FE	No	Yes	No	Yes
Observations	2040	2040	2039	2039
R-squared	0.797	0.953	0.839	0.993
Panel B. Patents				
female-inventor patents	9.525***	9.675***		
	(0.0720)	(0.0689)		
log female-inventor patents			0.861***	0.888***
			(0.00812)	(0.00747)
Field & vintage FE	No	Yes	No	Yes
Observations	1574	1574	1559	1559
R-squared	0.918	0.971	0.878	0.981

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Constant not reported.

Notes: Dependent variable is the number of male-authored articles (Panel A) / male-inventor patents (Panel B) by field and vintage, in levels (columns (1)–(2)) and in logs (columns (3)–(4)). The independent variable is the number of mostly-female-authored articles (Panel A) / mostly-female-invented patents (Panel B), in the corresponding levels or logs. Conventional standard errors are reported in all columns.

Table 7: First stage regressions

	Articles			Patents		
	(1) All	(2) Female	(3) Male	(4) All	(5) Female	(6) Male
Instrument: log corpus size	-1.5436***	-1.8491***	-1.4762***	-1.2858***	-1.3874***	-1.2791***
	(0.0131)	(0.0268)	(0.0148)	(0.0374)	(0.0534)	(0.0377)
Female-authored (df)	0.0012	0.0613***	-0.0472***	-0.0384***	0.0053	-0.0452***
	(0.0014)	(0.0027)	(0.0038)	(0.0037)	(0.0051)	(0.0035)
Observations	10,504,561	3,435,811	9,455,094	31,814,328	4,436,391	30,300,665

Standard errors in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Constant not reported.

Notes: The table reports the first-stage regressions for the 2SLS estimates in Table 8. The dependent variable is the log within-inside-good share $\ln(s_{j|\text{inside}})$, the endogenous regressor in the demand model, and the reported coefficient is on the excluded instrument, log corpus size, defined as the log number of innovations in the market-year. Each regression includes the female-authorship control *df* and absorbs cited-work vintage fixed effects, exactly as in the second stage. Estimation is by OLS on Frisch–Waugh–Lovell residualized variables, with standard errors clustered by citing year. Samples, columns, citer-group definitions, and the observation unit (cited work × citing-year) are identical to those in Table 8. Columns 1–3 report results for articles and columns 4–6 for patents. Within each domain, the columns correspond to all, female, and male citers. The article sample covers citing years from 1980 onward; the patent sample covers citing years 1990–2020, cited years 1976–2020, and forward citations only.

Table 8: Nested logit demand models

	Articles			Patents		
	(1) All	(2) Female	(3) Male	(4) All	(5) Female	(6) Male
σ (log inside share)	0.5922***	0.4570***	0.6271***	0.3764***	0.4667***	0.3844***
	(0.0119)	(0.0091)	(0.0135)	(0.0260)	(0.0227)	(0.0258)
Female-authored (df)	0.0007	0.0354***	-0.0169***	-0.0233***	0.0088**	-0.0281***
	(0.0005)	(0.0021)	(0.0021)	(0.0015)	(0.0030)	(0.0013)
First-stage F (KP)	13858.3	4755.7	9940.2	1179.4	674.6	1150.6
Observations	10,504,561	3,435,811	9,455,094	31,814,328	4,436,391	30,300,665

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$,*** $p < 0.001$. Constant not reported.
Notes: The table reports two-stage least squares (2SLS) estimates of a one-level nested-logit demand model for citations in which the inside good comprises all innovations — every field pooled into a single nest — and the outside good is non-citation. The dependent variable is the log mean-utility contrast $\ln(s_j) - \ln(s_0)$, where s_j is innovation j's share of the potential citation market (normalized to 100 potential citations per innovation) and s_0 is the outside-good share. The reported coefficient σ is on the log within-inside-good share $\ln(s_{j|inside})$, instrumented by log corpus size (the log number of innovations in the market-year). All columns control include cited-work vintage fixed effects (article publication year; patent application/grant year) and field fixed effects. Estimation is by Frisch–Waugh–Lovell residualization on the absorbed effects followed by 2SLS (ivreg2); standard errors (in parentheses) are clustered on the citing year, with the small-sample correction adjusted for the absorbed fixed effects. First-stage F is the Kleibergen–Paap rk Wald F. Columns 1–3 use the article citation panel and 4–6 the patent panel; within each domain the three columns differ by the gender of the citing author/inventor — All uses all citations, Female only citations made by (majority-)female citers, Male only citations by male/non-female citers — so shares and utilities are built from that citer group's citations. An observation is a cited work $\times$ citing-year cell. The article sample covers citing years 1980 onward; the patent sample covers citing years 1990–2020 with cited (application) years 1976–2020 and forward citations only (citing year cited year), restricted to WIPO technology fields. The Female columns have fewer observations because they retain only cited works receiving 1 citation from a female citer that year.

A Appendix

A.1 The range of possible equilibrium welfare effects

To see the range for the welfare effect, consider male and female innovation, with cumulative quality schedules $D^m(N^m)$ and $D^f(N^f)$, where under our assumptions both schedules are increasing and concave. The status quo numbers of innovations are N_0^m and N_0^f, so the status quo is $D_0 = D^m(N_0^m) + D^f(N_0^f)$, and innovator surplus is $IS(D) = M \ln\left(1 + D_0^{1-\sigma}\right)$. Because the $IS(D)$ function is strictly increasing in D, comparisons of consumer surplus are equivalent to comparisons of the term D.

The elimination of the female influx reduces D_0^f to D_1^f. After the female influx is removed, and holding male entry fixed, define $D_{\text{no}} \equiv D^m(N_0^m) + D_1^f$ where, of course, $D_{\text{no}} < D_0$. Now allow male entry to adjust to its counterfactual level N_1^m, and define $D_1 = D_1^m + D_1^f$. Equilibrium requires that the expected citations to the marginal male entrant are restored to their status quo level. In general, the expected citation return to a marginal male innovation is proportional to

$$K^m(N^m, D) = D^{m\prime}(N^m)\frac{D^{-\sigma}}{1 + D^{1-\sigma}}. \tag{9}$$

Thus, the status quo marginal return is

$$K_0^m = D^{m\prime}(N_0^m)\frac{D_0^{-\sigma}}{1 + D_0^{1-\sigma}}, \tag{10}$$

and the replacement-entry equilibrium satisfies

$$D^{m\prime}(N_1^m)\frac{D_1^{-\sigma}}{1 + D_1^{1-\sigma}} = K_0^m. \tag{11}$$

Given this, counterfactual IS can be as high as status quo IS. Consequently, the percentage change in IS from the female influx is not guaranteed to be positive.

Suppose instead that $D_1 > D_0$. Because $N_1^m \geq N_0^m$ and $D^m(\cdot)$ is weakly concave, $D^{m\prime}(N_1^m) \leq D^{m\prime}(N_0^m)$. Also, the term

$$h(D) = \frac{D^{-\sigma}}{1 + D^{1-\sigma}}$$

is strictly decreasing in D for $0 \leq \sigma < 1$. Therefore $D_1 > D_0$ implies $h(D_1) < h(D_0)$. Combining these two inequalities gives $D^{m\prime}(N_1^m)h(D_1) < D^{m\prime}(N_0^m)h(D_0)$, which contradicts the equilibrium condition requiring the marginal male return in the counterfactual to equal its status quo value. Hence, $D_1 \leq D_0$, or that equilibrium innovator surplus absent the female influx cannot exceed that in the status quo.

At the same time, because male entry weakly expands and $D^m(\cdot)$ is increasing, $D^m(N_1^m) \geq D^m(N_0^m)$. Hence, $D_1 = D^m(N_1^m) + D_1^f \geq D^m(N_0^m) + D_1^f = D_{\text{no}}$ so that total quality with

replacement satisfies $D_{\text{no}} \leq D_1 \leq D_0$. Since $IS(D)$ is increasing in D, it follows that

$$IS(D_{\text{no}}) \leq IS(D_1) \leq IS(D_0).$$

Without male replacement entry, IS rises. With male equilibrium entry, the welfare gain from the influx can be as low as zero.