Are Patents with Female Inventors Under-Cited? Evidence from Text Estimation[†]

Yael V. Hochberg
Rice University & NBER

Ali Kakhbod
UC Berkeley

Peiyao Li
UC Berkeley

Kunal Sachdeva
University of Michigan

January 2026

Abstract

Utilizing leading machine learning techniques to analyze the textual content and quality of patents, we demonstrate that patents with female lead inventors are under-cited relative to what would be expected had the lead inventor been male. Male inventors are the greatest contributors to the undercitation of patents with female inventors, followed by female inventors and male examiners, while female patent examiners appear to be even-handed. Using market reactions to patents suggests no average difference in market value by the inventor's gender. The results have potential implications for research conclusions that rely on citation-based assessments of patent quality.

JEL classification: J16, J24, J71, O30, C13
Keywords: Innovation, Gender, Patent, Machine Learning, Big Data, Inference

[†]Philipp Schnabl was the editor for this article. We thank Tania Babina, Daniel Bergstresser, Po-Hsuan Hsu, Jessica Jeffers, Amir Kermani, Leonid Kogan, Alan Kwan, Martin Lettau, David Levine, Zack Liu, Tim Liu, Dmitry Livdan, Ulrike Malmendier, Asaf Manela, Danqing Mei, S. Lakshmi Naaraayanan, Matthias Qian, Vira Semenova, Amit Seru, Carolyn Stein, Gauri Subramani, Dong Yang, Tengfei Zhang, and workshop and conference participants and discussants at UC Berkeley, NBER Productivity Innovation and Entrepreneurship Meeting, Midwest Finance Association (MFA), York University (Schulich), Mid-Atlantic Research Conference in Finance (MARC), Northeastern University (D'Amore-McKim), Boston University (TPRI IP Day), Rice University (Jones), University of Virginia (Darden), ETH Zürich, University of Rochester (Simon), HEC Paris Entrepreneurship Workshop, University of Colorado (Leeds) Winter Summit, The Financial Intermediation Research Society (FIRS), University of Michigan (Ross), and UMass Amherst (Isenberg) for their helpful comments and suggestions. This paper was previously circulated under the title "Inventor Gender and Patent Undercitation: Evidence from Causal Text Estimation."

E-mail: Hochberg (hochberg@rice.edu), Kakhbod (akakhbod@berkeley.edu), Li (ojhfklsjhl@berkeley.edu), Sachdeva (ksach@umich.edu).

Female inventors appear to face significant obstacles when seeking patents. Women face disparities in the patent approval process (Jensen, Kovács, and Sorenson (2018)), may face a higher bar for patent grants (Gavrilova and Juranek (2021)), and are underrepresented among inventors in patent applications and more generally (Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019); Hunt (2016); Reshef, Aneja, and Subramani (2021); Toole, Saksena, deGrazia, Black, Lissoni, Miguelez, and Tarasconi (2020)). As a result, assuming a patent is applied for and granted, we might well expect to observe a selection effect, whereby the average patent with a female inventor may be of higher quality than the average male inventor patent, and thus, on average, achieve a higher forward citation count. Perhaps surprisingly, however, a simple analysis of patent citations suggests that patents with female inventors receive *fewer* forward citations relative to male inventor patents (Jensen et al. (2018)), suggesting either that patents granted to female inventors are of lower quality, or, more concernedly, that the quality of their patents is not fully recognized in the form of forward citations, a commonly used proxy for patent quality.

Assessing whether a patent is undercited relative to its actual quality is not a trivial undertaking. Typically, citations serve as the de facto measure of a patent's quality, even though the measure is noisy. To determine whether female inventors face systematic obstacles to citations of their work versus simply producing lower-quality patents, the econometrician must disentangle actual quality from the citation outcome. In an ideal setting, the econometrician would either randomize underlying quality across genders or gender across patents. Natural experiments that mimic this ideal, or suitable instrumental variables, however, have been elusive.

In this paper, we utilize novel machine learning techniques that allow for the measurement of the contribution of gender to the citation of a particular patent based on its content and forward and backward similarity to the existing and future corpus of patents (Kelly, Papanikolaou, Seru, and Taddy (2021)). Our methodology builds on a burgeoning set of research in the computer science literature that attempts to identify gaps associated with binary variables using textual analysis (Khetan, Ramnani, Anand, Sengupta, and Fano (2022); Shao, Li, Gu, Qian, and Zhou (2021)). Our empirical methodology, which we label as Inference from Textual Analysis, or I-TEXT, utilizes neural network analysis of observational text data–adjusting for confounding features of the text, and accounting for patent novelty and impact–to infer the expected quality (number of forward citations) for a given patent under different lead inventor genders.[1] The methodology generalizes approaches from the computer science literature such as Veitch, Sridhar, and Blei (2020) and allows for various text encoder architectures (here, Longformer and SciBERT) to produce numerical "embeddings"—relatively low-dimensional vector representations of the contents of the patent—that, combined with measures of patent novelty (backward similarity to existing corpus of patents) and impact (forward similarity to the subsequent corpus of patents), preserve sufficient information for characterizing overall patent quality. Our goal is to reduce dimensionality yet preserve aspects of the patent content and its significance that can explain both the main independent variable of interest (gender) and the outcome (citations) while disposing of linguistically irrelevant information. The methodology also incorporates a gender propensity model to ensure a patent's text is not readily identifiable as male- or female lead inventor from the writing alone, in order to ensure a quality estimate conditioning on the opposite gender can be computed.

We use the resulting embeddings to fit two neural networks—one per gender of the lead patent inventor—using the embedding's numerical representations of the patents and the patents' novelty and impact as inputs, and forward citations as outputs. Each neural network represents a mapping from this data to citation counts. The first mapping is estimated using the subset of data where the patent has a female lead inventor, while the second mapping uses the data where the patent has a male lead inventor. Unlike the standard OLS approach, the neural network approach captures complex and often nonlinear relationships between inputs and outputs.

Having obtained parameters for each gender's citation-estimation model, we then take the sample of patent data and run it through both citation models. This produces two sets of estimated citation counts for each patent, holding the content of the patent constant and changing only the gender of the inventors. Our analysis employs the difference between the actual forward citations a patent receives and the estimated citations for the same exact patent had it had a male lead inventor instead.

Our main sample covers all utility patents granted by the U.S. Patent Office (USPTO) from 1976 through 2015. For our main analyses, we focus on the first inventor's name on the patent

[1] Throughout the paper we use the terms first inventor and lead inventor interchangeably to refer to the inventor whose name appears first in the list of inventors on a patent. Unless otherwise specified, when we refer to the gender of an inventor, we are referring to the lead (first) inventor.

(the "lead inventor"), and label patents as female lead inventor if the first inventor listed on the patent is female, and male lead inventor if the first inventor listed on the patent is male. Our results are robust to other labeling approaches of female versus male inventor patents, including restricting to single-inventor patents or majority-gender teams. Because the modal patent receives zero citations, our main analysis focuses on the subsample of patents that receive at least one citation.[2]

We begin by documenting that, consistent with Jensen et al. (2018), and controlling for a variety of patent characteristics, patents with female lead inventors receive fewer citations than those with male lead inventors, particularly when such patents are of high significance (as defined in Kelly, Papanikolaou, Seru, and Taddy (2021)). We then apply the I-TEXT methodology to obtain estimated forward citations for each patent under each lead inventor's gender.

Next, for each individual patent, we calculate our primary outcome variable, *Delta*, which we define as the difference between the *actual* number of forward citations received and the estimated forward citations had the inventor(s) been male. We use *Delta* as the dependent variable in simple OLS regressions that allow us to examine the heterogeneity in citation gender gaps, and include a variety of controls and fixed effects (such as for art unit and patent examiner) for added robustness.

Our baseline estimates suggest that patents with a female lead inventor would have received more citations if their lead inventor had instead been male, holding patent content fixed. The average gender gap for the sample, absent any controls, is -2.12. In other words, for any given patent, our trained models would estimate 2.12 fewer forward citations if the first inventor listed was female than it would for the exact same patent if the first inventor listed had instead been male. Estimating OLS models that further control for patent significance, art unit, examiner, and attorney fixed effects, we find that patents with a female lead inventor receive approximately 2.5 fewer citations than the same patent would receive had the lead inventor been male—approximately 13% fewer citations as compared to the mean for female lead inventor patents.

Identifying the *source* of this gender citation gap is clearly of interest. Undercitation of female lead inventor patents could be driven by the citations inventors themselves include in their patent applications, or by citations that are added by patent examiners. This in turn

[2]Analysis using estimations from the full dataset into the neural net is presented in the Online Appendix Table IA1, with coefficients similar to the intensive margin sample we use in the main analyses.

may depend on the gender of the examiner or inventor who could be adding the citation. Controlling for art unit, issue year, examiner, attorney, and assignee, we find that both patents with male lead inventors and patents with female lead inventors significantly undercite past patents with a female lead inventor, though this gap is approximately triple in magnitude for male lead inventors. When we look at the citations added by patent examiners, we find that male patent examiners mildly undercite patents with female lead inventors, while female patent examiners do not appear to undercite female lead inventor patents.

The results are robust to various alternative specifications and do not appear to be attributable to sample selection or model overfitting. The results are also robust to various approaches to defining a "female inventor" patent. For example, we obtain similar results when comparing patents with a single female inventor to those with a single male inventor or patents with inventor teams composed of a majority of female inventors versus patents with a majority of male inventors. Our results hold across Cooperative Patent Classification (CPC) major categories and subcategories of patent technology, with some heterogeneity by subcategories. Significant patents with female lead inventors in the Human Necessities CPC category are particularly undercited. We observe similar patterns when using the National Bureau of Economics (NBER) classification system. Undercitation of significant female lead inventor patents also appears to be particularly large in emerging technology fields.

Finally, we explore whether the public market's assessed economic value of patents varies with gender of the lead inventor. We use the economic value of a patent as measured by public markets from Kogan, Papanikolaou, Seru, and Stoffman (2017), which is generally considered forward-looking. We regress an indicator for female lead inventors against the real dollar value of the patent and find that the undervaluation of female patents disappears when we saturate the model with a standard set of fixed effects. Importantly, even when controlling for actual forward citations, estimated forward citations if the lead inventor had been male, we find that the gender gap is small and statistically indistinguishable from zero. The estimates suggest that public markets do not assign different values to male and female lead inventor patents, controlling for observables, despite the gap in citations for such patents. Put differently, expected economic value, as measured by market reactions, may be a more accurate proxy for patent quality than standard measures of realized forward citations.

Although our findings suggest that female inventors receive fewer citations to their patents than they would have for the exact same patent had they been male, further research will be necessary to establish *why* patents with female lead inventors are undercited. While the evidence we present strongly points to the presence of a systematic underrecognition of female lead inventor patents, the fact that market-based valuations do not differ materially by inventor gender once textual and contextual features are accounted for would be inconsistent with female lead inventor patents simply being of lower economic value. Moreover, the finding that the undercitation gap widens for highly significant patents suggests a substantial–and arguably concerning—-discrepancy between the credit such patents receive and the credit they would receive if their lead inventor were male. While these patterns are consistent with some form of discrimination or bias, they alone cannot distinguish between taste-based discrimination (citing behavior driven by preference rather than expected quality), statistical discrimination (citers inferring lower quality from female names), or other mechanisms, such as differences in network structures or visibility.

Critically, the data cannot speak directly to the intent of the citers (whether examiners or inventors) or conclusively identify the precise channel(s) behind the undercitation of female-led work. For instance, male inventors appear to be the biggest contributors to the gender gap in forward citations, but female inventors also undercite female-led patents to some degree. Such patterns could stem from homophily in professional networks, differences in information flow, or entrenched norms that treat female inventor inventions less prominently-—all of which might masquerade as discrimination. Hence, while the results demonstrate that female inventors receive fewer citations than their male counterparts for patents of similar textual quality and context, they do not establish definitively whether these behaviors constitute taste-based or statistical discrimination, nor do they fully rule out alternative social or structural explanations. We note, however, that our findings regarding the similar market value of patents for both genders suggests statistical discrimination is unlikely to be the mechanism behind the undercitation pattern we uncover.

Our primary contribution is thus to provide a new, measurement-focused analysis of how gender-based undercitation manifests in patent data. By employing recent advances in natural language processing, we introduce a text-embedding approach that refines the measurement of patent quality and allows us to document a systematic gap in forward citations for patents with female inventors relative to those with male inventors.

Our findings have potentially important implications. First, the literature has highlighted that innovation is motivated by the expected profits derived from the property rights granted to patentees, Moser (2005, 2013).[3] If women are not equally recognized for equivalent patents, this may discourage them from entering the innovation economy, potentially reducing contributions from half of the population, and exacerbating the already substantial wedge between men and women in science, technology, engineering, and mathematics (STEM) fields (Beede, Julian, Langdon, McKittrick, Khan, and Doms (2011), National Science Foundation (2013)), leading to further inefficient allocation of labor. Second, our findings raise concerns regarding the validity of research that relies on forward citations of patents as a measure of patent quality. To the extent that patents with female inventors are systematically undercited relative to their male counterparts, using forward citations as a measure or control for quality may be contraindicated. Given the large literature that relies on forward citations, a re-examination of prior findings may be warranted.

Our findings reveal that female-led patents are systematically undercited relative to comparable male-led patents, contributing to a broader literature on barriers faced by women in knowledge creation and dissemination. Recent studies document that women may receive less credit for their academic work than men. For example, Sherman and Tookes (2022) show that, controlling for research productivity, female finance scholars tend to be at lower-ranked institutions, earn less, and are less likely to become full professors. Sarsons, Gërxhani, Reuben, and Schram (2021) and Sarsons (2017) find that female economists get less recognition for coauthored work. Koffi (2021) highlights that female economists' top-tier papers are cited less in elite journals—especially by men. Chawla (2016); Koffi and Marx (2023) identify parallel gender gaps in other scientific fields. In contemporaneous work to ours, Koffi (2024) shows that undercitation disproportionately affects women's papers, and Iaria, Schwarz, and Waldinger (2024) report that despite reduced hiring gaps over time, women still receive fewer citations and promotions. Hengel and Moon (2020) document that female-authored papers in top economics journals are on average of higher quality than male-authored ones.[4] By uncovering analogous gender-based citation disparities in patents, we show that these inequities in recognition extend beyond academic publishing into intellectual property. Like much of this

[3]In related research, the marginal investor values patents Aghion, Van Reenen, and Zingales (2013); Hall, Jaffe, and Trajtenberg (2005); Hirschey and Richardson (2004); Hirshleifer, Hsu, and Li (2013).

[4]Not all research finds negative gaps for women; Card, DellaVigna, Funk, and Iriberri (2020) conclude that top economics journals actually cite male-authored papers at lower rates.

work, however, we cannot definitively pinpoint the driving factors—whether discrimination, preferences, networks, or other mechanisms.

Finally, our methodological contribution lies in demonstrating the utility of text-embedding techniques for capturing nuanced dimensions of patent quality and content. Building on a rapidly growing branch of "big data" approaches in economics and finance (e.g., Gentzkow, Kelly, and Taddy (2019a)),[5] we offer a new way to systematically compare the content of patent documents and thereby better isolate the extent to which undercitation may reflect gender bias. This measurement strategy is particularly relevant for future work, as it provides a more granular framework to investigate how textual similarity, rather than merely technology class or keyword matching, can aid in understanding variations in patent citations and outcomes.

1 Patent Data

Our patent data comes from the USPTO.

1.1 Gender of Inventors

Our main independent variable of interest is the gender of the lead inventor of a patent (equivalent to the first author on a research paper). The person named first on a patent is typically the primary contributor.[6] The first name listed on the patent may be particularly salient, in a manner similar to how academic papers with multiple authors are often referred to as *"FirstAuthor et al."* We obtain the gender of the inventors on the patent from PatentView.

1.2 Patent Citations

Patents' forward citation counts are obtained using data from the USPTO. While forward citations have historically been used as a proxy for patent quality, the key point of our analysis is to determine whether this measure is systematically biased downward for patents with female

[5]For examples, see Athey and Imbens (2019); Bellstam, Bhagat, and Cookson (2021); Cong, Liang, and Zhang (2019); Erel, Stern, Tan, and Weisbach (2021); Gentzkow et al. (2019a); Gentzkow, Shapiro, and Taddy (2019b); Hanley and Hoberg (2019); Hansen, McMahon, and Prat (2018); Li, Mai, Shen, and Yan (2021); Loughran and McDonald (2016); Rouen, Sachdeva, and Yoon (2022); Routledge, Sacchetto, and Smith (2017).

[6]In further robustness, we consider single-inventor patents and the gender of the entire team.

inventors. We therefore distinguish between patent forward citations (the easily observed outcome for a patent) and the *quality* of the patent, which is the measure of interest and can be mediated by using the content of the patent in addition to its novelty and impact.

1.3 Examiner Versus Inventor Added Citations

Typically, patent applications include a list of related patents and supporting material. Citations to patents may be added in two ways. First, inventors cite precedent patents in their applications. Second, examiners will identify additional citations they deem "missing" from the patent, and request that these be included (Farre-Mensa, Liu, and Nickerson (2022)). Starting in 2001, and more clearly since 2003, the USPTO discloses whether the citation originated from the examiner or the inventor. For the purposes of the analysis, studying the source of a citation, we create additional citation counts that separate the citations that examiners and inventors added.

We consider the gender of the examiner and the propensity to cite each gender in part of our analysis. One problem with this, however, is that the USPTO patent data does not disclose the gender of examiners. As a result, we infer gender from third-party sources, (Graham, Marco, and Miller (2018)). To disambiguate the gender of the examiner, we implement a name disambiguation algorithm similar to that of Desai (2019). We use the first name of the lead examiner to identify their gender (Tzioumis (2018)).

Starting with the PatentView data, we obtain the first names of each patent examiner. Due to their prominence, we rely on the name of the lead (first) examiner for patents with multiple examiners. Next, we classify the gender of the patent examiner using state-level data on the frequency of names obtained from the Social Security Administration (SSA) (Comenetz (2016)).[7] We classify a name as belonging to a particular gender when it is associated with that gender in more than 99% of observations. If the first name does not match the SSA dataset, our second step uses a similar process but utilizes a cross-country dataset from the World Intellectual Property Organization (WIPO) (Martinez, Raffo, Saito, et al. (2016)).

[7] Gender of inventors are obtained from PatentView dataset.

1.4 Patent Classification

Our main sample covers all utility patents granted by the U.S. Patent Office (USPTO) from 1976 through 2015. Our sample of patents comes from the USPTO's Patent Examination Research Database dataset.

When an inventor files a patent application with the USPTO, the application is assigned a USPC class and subclass based on its technology field. The application is then assigned to an "art unit" comprised of several examiners specializing in that particular technology class and subclass. We use the art unit to which the patent is assigned as our proxy for technology-type grouping. Our baseline sample contains 735 art units and 15,770 patent examiners. As an alternative to the art unit, we employ the Cooperative Patent Classification (CPC) of the patents and the NBER patent category, which is also reported in the USPTO PatentView database.[8]

1.5 Attorneys and Assignees

Patents are typically filed with the assistance of a patent attorney, who may file many of them on behalf of different inventors. The USPTO also reports the initial assignee of the patent. We use these identifiers as they help us account for possible commonalities in writing style across patent attorneys and firms that may influence the textual content of the final submission.

1.6 Patent Significance

While the patent text is crucial for measuring quality, *when* it was written is equally important. The same exact patent filed in 2010 versus 2020 may have different significance, as the corpus of precedent and subsequent patents will differ. For example, a patent on artificial intelligence algorithms written and filed in 2015 will likely have a higher potential to be distinct from prior patents (novelty) than the same patent if filed instead in 2024, while also being more likely to be influential on future patents (impact). We therefore also consider information about the importance of the focal patent relative to the patents that appear both before and after, utilizing measures of patent novelty and impact as in Kelly et al. (2021).

[8] Note, the NBER patent categories are truncated at the end of our sample.

1.6.1 Measuring Contextual Importance

To account for the context of when the focal patent was written, Kelly et al. (2021) consider the similarity of the text to the text of patents that appear either before or after it is granted. Specifically, they use the Term Frequency-Backward Inverse Document Frequency (TF-BIDF) vectors from Kelly et al. (2021). Here, the TF-BIDF is defined as a vector where each entry corresponds to a unique word in patents. TF, BIDF, and TF-BIDF are defined separately as follows:

$$\text{TF}_{pw} = \frac{c_{pw}}{\sum_k c_{pk}},$$

where c_{pw} is the number of times the word w appears in patent p (similar to c_{pk} for word k).

$$\text{BIDF}_{tw} = \log\left(\frac{\text{\#patents prior to } t}{1 + \text{\#documents prior to } t \text{ that include term } w}\right),$$

and for any pair of patents i and j, the TF-BIDF for patent i is a vector of length $|W|$ where W is the union of unique words in patents i and j: for any $w \in W$

$$\text{TF-BIDF}_{w,i,t} = \text{TF}_{iw} \times \text{BIDF}_{tw},$$

where $t \equiv \min(\text{filing year of i}, \text{filing year of j})$.

Given the TF-BIDF vectors of patents i and j, the researcher can compute the similarity between the two as

$$\text{sim}(i,j) = \frac{\text{TF-BIDF}_{i,t}}{||\text{TF-BIDF}_{i,t}||} \cdot \frac{\text{TF-BIDF}_{j,t}}{||\text{TF-BIDF}_{j,t}||}.$$

1.6.2 Measuring Impact

To measure the impact of the patent, Kelly et al. (2021) then calculate the similarity between the focal patent and patents granted in the first five years thereafter. To do this, they measure:

$$F_i^5 = \frac{\sum_{j \in S_F^5} \text{sim}(i,j)}{|S_F^5|}, \tag{1}$$

where S_F^5 is the set of patents issued within 5 years after the issue year of patent i, and

F_i^5 measures the *Forward Similarity* of a focal patent, with greater values indicating more impactful patents. They denote $\text{sim}(i, j)$ to be the cosine similarity between the TF-BIDF vectors of the two patents i and j. Intuitively, if future patents use language similar to that of a focal patent, it indicates that the focal patent is likely impactful.

1.6.3 Measuring Novelty

It is also important to measure how similar focal patents are to patents in the preceding five years. To do this, Kelly et al. (2021) measure the backward similarity using:

$$B_i^5 = \frac{\sum_{j \in S_B^5} \text{sim}(i, j)}{|S_B^5|}, \tag{2}$$

where S_B^5 is the set of patents issued within 5 years before the issue year of patent i, and B_i^5 measures the *Backward Similarity* of a focal patent. Taking the inverse of the backward similarity $Novelty = \frac{1}{B_i^5}$ creates a novelty measure. Intuitively, patents with smaller B_i^5 (larger Novelty) values are associated with being more novel, as their text is distinct from what appeared in the prior five years.

1.6.4 Measuring Significant Patents

Following Kelly, Papanikolaou, Seru, and Taddy (2021), to compute the overall significance of a patent, we scale the forward similarity (impact) by the backward similarity (novelty) of this patent. This is defined as:

$$\textbf{Significance}_i = \frac{F_i^5}{B_i^5}, \tag{3}$$

where greater values indicate more significant patents. For ease of interpretation, we further scale the Significance variable to have a mean of zero and a unit standard deviation.

Descriptive statistics for our sample are presented in Table 1.

2 Inference from Text Analysis (I-TEXT)

Our desired analysis presents both methodological and computational challenges. First, we must represent complex and often subtle differences in the text of the patents in a parsimonious and computationally useful form. Second, we need to relate that text to forward citations. Finally, we must estimate the number of citations based on the gender of the inventor.

Below, we outline our empirical strategy. First, we discuss how we create a low-dimensional representation of text that encapsulates the information necessary to distinguish patent quality. Second, using this representation, we provide an overview of the I-TEXT methodology and how we train our model. Finally, we discuss the key identification assumptions implicit in our approach and their validity.

2.1 Text Embedding

We rely on the Bidirectional Encoder Representations from Transformers (BERT) family of models, which transform each piece of text into a low-dimensional numerical vector. Developed by Google (Devlin, Chang, Lee, and Toutanova (2018)), BERT has become the leading approach in many commercial applications (we refer the reader to the internet appendix B for a detailed discussion of neural networks and BERT). BERT uses the attention mechanism to construct embedding vectors that are numerical representations of the text, which preserve both the meaning of individual words and the underlying context of each word (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, and Polosukhin (2017)).[9]

For our main analysis, we use an encoder, Longformer, which is a special version of a "Robustly Optimized BERT Pretraining Approach" (RoBERTa) adopted to excel at representing long texts using newly developed mechanisms like global and local attention (Beltagy, Peters, and Cohan (2020)) (See appendix subsection C.1 for details on our encoder architecture).[10] In our application, the encoder model produces a low-dimensional representation

[9]See Jha, Liu, and Manela (2025) for an excellent discussion of BERT.

[10]In contrast to a traditional BERT model, RoBERTa is only trained on masked token prediction, is trained using a larger sample that contains long texts, and is trained over a longer period of time. These differences make RoBERTa a popular encoder model that outperforms BERT in many applications (Liu, Ott, Goyal, Du, Joshi, Chen, Levy, Lewis, Zettlemoyer, and Stoyanov (2019)).

with 768-dimensional embeddings to represent the text of the patent.[11]

To ensure the robustness of our estimates and their independence from specific text representations, we conduct an additional test using SciBERT, fine-tuned on scientific writing. Our results demonstrate that these alternative embedding models yield highly consistent results, reinforcing the validity of our findings.

2.2 I-TEXT Methodology

Having created low-dimensional representations of patent text, the second challenge is establishing the relationship between this data and patent forward citations. To do so, we introduce a novel machine learning technique, which we label as "Inference from Text Analysis (*I-TEXT*)." that allows us to estimate the gap in outcome (forward citation counts) between two groups denoted by a binary independent variable (gender), using the text of the patent and other inputs as a mediator for quality.[12] I-TEXT builds on recent advances in computer science, including Khetan et al. (2022); Shao et al. (2021); Veitch et al. (2020). Inference text analysis allows us to use the text of patents as a mediator to identify the role of gender for patent citations. To the best of our knowledge, ours is among the first papers in economics to apply deep learning to inference from language.

I-TEXT is a neural network-based architecture that produces estimates under two conditions (male versus female inventor in our case). As shown in Figure IC2, in our analysis, the input data for fitting the neural networks (commonly referred to as X in computer science) contains four types of information: the text of the patents, impact and novelty, gender indicators of the inventor(s), and the observed number of forward citations to the patents. There are four neural networks that need to be fitted: a Longformer model for generating text embeddings, a logit-linear model that maps embeddings, novelty, and impact to propensities, and two one-hidden-layer 200-node ReLU-activated neural networks that map from embeddings, in conjunction with novelty and impact measures, to the estimated number of citations for male and female inventors, respectively. The final loss function is a weighted average of the losses of these four neural networks.

The I-TEXT methodology has three key steps (See Figure 1). First, it uses a language model

[11] Patentview provides the full text of the patent. Importantly, this text does not contain header information such as title, author, firm, and location, and the header information is not an input to our model.

[12] See subsection C.2 for a more in-depth discussion.

to encode the text of the patent. As discussed above, our main approach uses Longformer, a transformer-based model, to transform the text of each patent into a low-dimensional numerical vector. The embedding vectors are numerical representations of the text that look to preserve both the meaning of individual words and the underlying context of each word. Second, we obtain the impact (forward similarity) and novelty (backward similarity) of the patent. Third, we feed the embedding vector, patent novelty, and patent impact into the neural networks that estimate forward citations by gender.

I-TEXT estimates the number of citations an inventor would have received if that person were assigned the opposite gender. This is accomplished by fitting two neural networks, where each model represents a mapping from embedding vectors, patent novelty, and patent impact to our outcome variable, forward citations, with the first mapping fitted using the subset of patents with a female lead inventor and the second mapping fitted using the subset of patents with a male lead inventor. The two estimated mappings, combined with the flexibility of neural networks, allow us to approximate the true mappings.

Armed with our two mappings, we can ask the following: how many citations would a patent whose lead inventor is female have received if the lead inventor had instead been male, and vice versa?

In parallel, regardless of the gender indicators, each embedding is passed through a propensity network that is used to estimate the propensity of that particular patent to have had a female lead inventor. We use this propensity score to identify patents that are very likely (97%+ probability) to have lead inventors of one gender or the other (irrespective of quality), which are then dropped from the sample (as discussed in section 1). Finally, the model's output is two sets of citations (one from male and one from female) and a set of female propensity.[13]

The framework can be expressed more formally in mathematical terms. We denote the text in the ith patent as W_i. We fine-tune the Longformer model $\phi(\cdot)$ to map W_i to its embedding of the patent text. We concatenate the embeddings with their novelty and impact scores to produce a vector Z_i. Then, we use a logit-linear network $\hat{g}$ to map Z_i to a real number, which

[13]The citation estimates using the true gender network are also saved and used for computing the gap in citation between the two genders.

represents the probability that this patent has a female lead inventor:

$$\hat{g}(Z_i) = P(T_i = 1|Z_i), \tag{4}$$

where T_i is the gender of the patent i's inventor (0 if male and 1 if female).

In addition, we have two citation networks $\hat{Q}_1$ and $\hat{Q}_0$. $\hat{Q}_1$ maps an embedding vector concatenated with novelty and impact scores to the estimated number of citations if the patent has a female lead inventor, and $\hat{Q}_0$ maps an embedding vector concatenated with novelty and impact scores to the estimated number of citations if the patent has a male lead inventor. Mathematically, we define a piecewise mapping $\hat{Q}$ that represents the two networks:

$$\hat{Q}(T_i, Z_i) = \mathbb{E}(Y_i(T_i)|Z_i), \tag{5}$$

where $Y_i(0)$ and $Y_i(1)$ denote the hypothetical outcomes of the ith patent. In our case, these hypothetical outcomes are the number of forward citations. Given these mappings represented by neural networks, for a set of N patents, we can then estimate the gap in citation between the two gender groups using the following equations:

$$\begin{aligned} \text{Gap} &= \frac{1}{N}\sum_{i=1}^{N}[\mathbb{E}(Y_i(1)|Z_i) - \mathbb{E}(Y_i(0)|Z_i)] \\ &= \frac{1}{N}\sum_{i=1}^{N}\left[\hat{Q}_i(1, Z_i) - \hat{Q}_i(0, Z_i)\right], \end{aligned} \tag{6}$$

While the model estimates the quality of the patent based on textual content, novelty, and impact, forward citations may also vary with observable characteristics of the patent that are not related to quality, such as the art unit or technology class, the attorney who filed the patent, or the assignee, which are also not included in the patent text. Next, we pass the difference between actual forward citations and estimated forward citations (our text quality measure) through OLS regressions with fixed effects for the above factors. Our observation for a patent then consists of the actual forward citations for the patent, the number that would be estimated if the lead inventor was female, the number that would be estimated if the lead inventor was male (text quality), and the actual gender of the leader inventor (1/0, where 1 is female, and 0 is male), and other patent characteristics.

2.3 I-TEXT's Identification Assumptions

While I-TEXT is a flexible and powerful framework, there are key assumptions that the econometrician must consider in interpreting its output. We discuss each of them below.

2.3.1 Identifiable Gap

The first necessary condition is that the text of the documents, when combined with any other inputs into the neural networks (here, impact and novelty), must make the gap identifiable. In other words, the gap being measured by the econometrician must be directly quantifiable from the text and the supplemental characteristics. This condition is similar to an exclusion restriction in standard identification strategies. Because this cannot be formally tested, the validity of this condition must be carefully examined and potentially falsified, considering alternative channels.

In the context of this paper, we wish to measure the relationship between patent content, novelty, impact, inventors' gender, and forward citations. Our identifying assumption is that the quality of the patent should be measurable by the content (text) of the patent and these significance measures. Patent examiners read the patent application's text, comparing it to the corpus of existing patents, in order to evaluate the proposed patent's quality before granting it. Further, the text of the patents, by construction, should contain all relevant information related to a patent and the invention it describes. As a result, this necessary condition is likely well satisfied in our context.

2.3.2 Embedding Method Extracts Semantically Meaningful Information

The second necessary condition is that the embedding method extracts semantically meaningful text information relevant to estimating the group indicator, T, and outcome, Y. In our setting, this means that embedding, a lower-dimensional representation of the text, is sufficient to capture the content of the text.

The pre-trained text encoder we use for our main analysis, Longformer, has been shown to excel at multiple linguistic tasks such as question answering and text classification. Beltagy et al. (2020) show that Longformer achieves over 75% accuracy on answering trivia questions and around 95% on text classifications. In addition, a large body of work has demonstrated the effectiveness of fine-tuning Longformer in text analysis. Li, Wehbe, Ahmad, Wang, and Luo (2022) shows that the fine-tuned version of Longformer performs well in clinical applications such as question answering and named entity recognition (identifying words or phrases that represent named entities such as proper nouns). Khandelwal (2021) shows that a fine-tuned Longformer model performs well in classifying posters' stances via their discussions about controversial issues on social media. The other encoder we consider as additional robustness analysis, SciBERT, has also been shown to be effective in text analysis (Beltagy, Lo, and Cohan (2019)).

2.3.3 Conditional Outcome and Propensity Score Models are Consistent

Our third and final necessary condition is that the conditional outcome and propensity score models are consistent, meaning that the two groups defined by our binary variable of interest should have common support. Within the context of this paper, the primary concern is that women and men can differ substantially in writing style, which may in some way translate into the text of their patent, and that such differences may lead to an inference of undercitation that is due to writing style rather than inventor gender.

To ensure common support, we follow the approach in Veitch et al. (2020) and estimate the propensity model using a one-layer logit-linear neural network, where the objective function is the binary cross-entropy between the estimated gender indicator and the true gender indicator. Using the patent text, this neural network's output is the estimated probability that a female lead inventor is the writer of the patent. We also filter on patents that the model identifies the gender with high confidence, less than 3% and greater than 97%.

2.4 Model Fit

We evaluate the model's fit and accuracy in detail in the Appendix D. When trained on 80% of the data and tested on the remaining 20%, our estimated text-based quality measure has a highly significant positive correlation with the actual number of citations: 90% fall within one standard deviation of actual values, and most of them are within 5 citations from the actual values. These results indicate that the model captures an economically meaningful share of the variation. The Appendix provides a full discussion of the optimization objective, training process, and diagnostic checks supporting these conclusions.

3 Do Citations to Patents Differ By Gender of Inventor?

Do forward citation counts for patents differ based on the gender of the lead inventor? We begin by examining the differences in actual forward citations between female and male lead inventor patents without model adjustment. Next, we calculate the I-TEXT implied gender gap based on patent text, patent novelty, and patent impact. We then further utilize the I-TEXT output to refine the estimates of differences in forward citation counts by gender using simple regression analysis to account for non-quality-related patent characteristics (e.g. art unit, year, etc.).

3.1 Difference Between Genders Without Model Adjustments

We begin by plotting the unconditional differences in citations by gender. The histogram for citations for male and female lead inventors is plotted in Panel A of Figure 2, and visually demonstrates that females receive fewer citations than males.[14]

On average, male lead inventors receive significantly more citations than female lead inventors (23.6 citations for males, 19.2 for females, F-stat = 715.6, Table 1). Testing the difference in distributions, we find a Kolmogorov-Smirnov statistic of $D = 0.062158$, with a p-value of $= 2.2 \times 10^{-16}$, further suggesting that male- and female-lead inventor patents' forward citation counts come from different distributions.

We more formally consider the contribution of lead inventor gender to patent forward citations by estimating the following OLS model:

$$Y_i = \beta_1 I\left(FemaleInventor_i\right) + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i, \quad (7)$$

where patent and year are represented by i and t, respectively. Y_i is our outcome of interest, actual forward citations. Our specification includes fixed effects for examiner ($\delta_{examiner}$), attorney ($\delta_{attorney}$), assignee ($\delta_{assignee}$), art unit ($\delta_{ArtUnit}$), and year of grant ($\delta_{GrantYear}$). All standard

[14]The natural logarithm transformation of these distributions is presented in Figure IA2.

errors in this paper, unless otherwise noted, are double-clustered by patent issue year and art unit. β_1 is our coefficient of interest; a positive value would indicate that female lead inventor patents receive more citations than males.

Panel A of Table 2 presents the estimates from the model for the sample of patents which received at least one forward citation (most patents do not receive any forward citations; in general, female lead-inventor patents are less likely to receive any citations at all). The estimates suggest that conditional on receiving citations, female lead inventor patents receive between 0.88 to 2.4 fewer citations than male lead inventor patents, depending on the controls included.

Given that the expected selection effect from prior literature might predict that we would see higher quality patents–and thus, higher citation counts–for female lead inventor patents, the patterns from this simple analysis raise an interesting set of possibilities: either the female lead inventor patents being approved are of lower quality than those from males, or, despite the higher bar for approval of patents with female inventors demonstrated by past studies, the quality of these patents is not being appropriately reflected in citation counts.

3.2 Gender-based citation gap from I-TEXT

The I-TEXT methodology allows us to compute the gender-based citation gap (gap for short), mediating for patent content, novelty, and impact. Using our neural network mappings, we calculate the gap for a patent with a female lead inventor, as defined by (Equation 6), to be -3.12. This means that patents with female lead inventors tend to have 3.12 fewer forward citations compared to what they would have been expected to receive had the lead inventor instead been male.

3.3 Computing Delta in Citations

While the gap provides us with the overall difference between the two gender groups, controlling for the content, novelty, and impact of the patent, a number of observable characteristics not included in the patent text may affect forward citations through non-quality-related channels that the I-TEXT embeddings would not pick up, such as the art unit, technological classification, or tendencies/additions of the specific patent examiner assigned to the patent.

To control for these characteristics, we utilize a regression framework in which the dependent variable is the difference between the number of actual forward citations to a patent and the model-implied citations if the patent had a male lead inventor, i.e. under the male lead inventor neural network.

Specifically, we calculate the difference between actual and estimated text quality as:

$$Delta_i = ForwardCitation_i - Forward\widehat{Citation}_{i|T_i=0}, \tag{8}$$

where $ForwardCitation_i$ is the actual number of citations to a given patent and $Forward\widehat{Citation}_i$ is our patent text quality measure which is computed as the number of citations implied by the I-TEXT model if the lead inventor had been male. The measure holds the reference gender constant (in this analysis, male) to establish a uniform reference point for evaluating all patents.[15]

The *Delta* measure captures the citation bias at the patent level, measuring the relative undercitation of a patent. Specifically, *Delta* is calculated as the difference between the actual number of citations a patent receives and the number of citations it would have been expected to receive if it had a male lead inventor, based on a citation estimation model. A positive *Delta* indicates that a patent has been overcited relative to this expected benchmark, while a negative *Delta* indicates undercitation. This measure isolates the extent to which a patent's citation count diverges from the estimated number, allowing us to assess patterns of citation bias. Importantly, any factor that affects estimated-if-male citation values should also affect actual citation values. *Delta* captures the relative difference in how the underlying driving factors affect actual versus estimated-if-male citation counts.

As such, $Delta_i$ is measured in *citations*. A *Delta* of -1 indicates that the patent received one fewer citation in actuality than estimated by our model had the lead inventor been male. A *Delta* of +1 indicates that the patent received one more citation than our model would estimate had the lead inventor been male. Thus, a negative coefficient on *Female* in regressions using $Delta_i$ as the dependent variable indicates that female-led patents are, on average, under-cited relative to what the model would suggest. Because $Delta_i$ can be positive or negative, its sign and magnitude are inherently contextual—-our estimates report average differences in this

[15] Note that while *Delta* is computed relative to the male model as a benchmark, we could alternatively conduct the same exercise comparing to the female implied model as a benchmark.

residualized outcome.

As an illustrative example, consider a particular patent that, when its text, novelty, and impact are input into a neural network calibrated to estimate its forward citation count assuming a male lead inventor, is estimated to receive 15 forward citations. In actuality, however, this patent has been cited only 12 times. The computed *Delta* for this patent is -3, reflecting that the actual citation count is lower than the model-expected count under a male lead inventor. Had the actual number of citations received instead been 17, *Delta* for the patent would then be +2—the *Delta* value is not necessarily negative, as can be seen in Figure 3.

Importantly, we do not assume that $ForwardCitation_i$ and $Forward\widehat{Citation}_{i|T_i=0}$ are independent—indeed, they should be correlated. The key object of interest is the residual component: *Delta*, which we interpret as the *citation premium*—how many more citations a patent receives relative to the expected citation count based on our model. By construction, $\widehat{ForwardCitation}_i$ absorbs the variation explained by the patent's content (topics, terminology, and other semantic/linguistic features the model captures), the patent's novelty, and the patent's impact. Thus, *Delta* captures realized citations beyond that text- and novelty- and impact-estimated baseline (and beyond the fixed effects). Importantly, since the expected number of citations is determined solely by text, novelty, and impact, changing the gender indicator does not alter the expected count. Therefore, any gender-related difference in *Delta* reflects variation in the actual number of citations received.

Before incorporating *Delta* into our main analysis, we first examine its time-series and cross-sectional properties by gender. Panel A of Figure IA5 plots the evolution of *Delta* over time for patents with male and female lead inventors. In this exercise, we demean *Delta* using each year's average. Throughout the sample period, patents with female lead inventors exhibit negative average *Delta*, indicating that in each year, the difference between forward citations and model-based expected citation counts are lower for females than males, on average. The magnitude of *Delta* varies considerably over time, with particularly pronounced differences during the 1980s and 1990s.

We also examine how *Delta* varies across technology classes by gender. Panel B of Figure IA5 plots the average *Delta* by CPC class and gender, after demeaning *Delta* within each technology class. Across nearly all fields, patents with female lead inventors exhibit negative average *Delta*. The magnitude of these differences, however, varies substantially across technology classes, suggesting that the gender gap in predicted versus realized citations is not uniform across fields.

3.4 Undercitation of Female Lead Inventors Using Delta

We then relate *Delta* to various patent attributes unrelated to quality, such as examiner and art unit, in a regression framework. To do so, we re-estimate models of the type described in Equation 7, replacing the actual number of forward citations for a patent with *Delta*. The estimates are presented in Panel B of Table 2. The estimates suggest that patents with female lead inventors are undercited relative to what would be expected had the patent content and context remained exactly the same, save for the gender of lead inventor being male: in the strictest specification, patents with female lead inventors receive nearly 2.5 fewer citations (11% of the unconditional sample mean) than would be estimated if they had had a male first inventor instead for the exact same patent, controlling for other characteristics and mediating for content, novelty, and impact using I-TEXT.[16] As expected, the magnitudes are relatively unchanged when including patent issue year, art unit, examiner, attorney, and assignee fixed effects. The relative stability in estimates suggests that our analysis does not suffer from a correlated omitted variable, e.g. Oster (2019). Overall, the estimates suggest that female lead inventors are undercited, on average, relative to what their same patents would have received if the first name on the patent had been that of a male inventor. These differences in forward citations cannot be explained by differences in art units, time trends, or differences in assignees or examiners.

4 Cross Sectional Heterogeneity

Next, we explore whether these patterns of undercitation are uniform across technology classes.

[16]For robustness, we compute the citation gap over the full sample using the model trained on this balanced subset, and we see the gap is significant with similar point estimates as shown in Table IA1.

4.1 Technology Class

Our main approach is to break down the technology grouping using the NBER subcategories. Specifically, we estimate the following:

$$\begin{aligned} Y_i &= \beta_1 I\left(FemaleInventor_i\right) + \beta_2 I\left(FemaleInventor_i\right) \times \left(Subcategory\right) \\ &\quad + \delta_{PatentSubcategory} + \delta_{GrantYear} + \delta_{Customer \times Examiner} + \varepsilon_i. \end{aligned} \tag{9}$$

Figure 4 presents the linear combination of the female lead indicator and the interaction coefficients (Female Lead Inventor × Subcategory) for this analysis graphically. The finer category classification exhibits somewhat more heterogeneity than the major classes. Importantly, in all specifications, we include patent subcategory fixed effects to account for the average level of citations in a given subcategory. As can be seen clearly in Figure 4, for the vast majority of the technology subcategories, the estimates suggest that patents with lead female inventors are cited significantly less than they would have been expected to have been with a male lead inventor instead, and these citation undercounts are often substantial in magnitude.

4.2 Established Versus Emerging Fields

An interesting question is whether the patterns we see across technology fields relate in some way to whether women are patenting in an established field versus in an emerging field of technology. It is possible that newer fields may not present as many barriers to entry or pre-existing biases for female inventors and researchers, given the lack of an established history of research and researchers, and that we may expect undercitation patterns to be larger or concentrated in more established fields. It is also possible that newer fields are smaller and more competitive and insular, with higher barriers of entry for female inventors. On the other hand, the underlying forces that lead to undercitation for patents with female first inventors may be unrelated to the nature of the field, and relate to gender norms or perceptions more generally, in which case we would not expect to see a difference.

We denote a patent to be from an "emerging field" if the art unit first appeared within five years of the patent being granted.[17] Using this definition, we re-estimate our baseline specifi-

[17] Importantly, we observe that emerging fields are not a function of older patents and emerge consistently throughout the time period of the sample, Figure IA3.

cation by including an indicator for emerging fields and its interaction with the indicator for a female lead inventor. Our primary focus is on the coefficient of the interaction term, which captures the relationship between female lead inventors and emerging fields.

The estimates are presented in Table 3. Panel A presents estimates where the dependent variable is actual forward citations, and Panel B presents estimates where the dependent variable is $Delta$. Notably, in Panel A, patents in newer fields appear to receive more citations overall, on average, than patents in existing fields, consistent with the evolution of the novelty of inventions. In addition, when moving to Panel B, our estimates show that the main result still holds—depending on specification, patents with female first inventors receive 2.2 to 2.45 fewer citations than would be expected had the lead inventor instead been male. Moreover, the estimates also suggest that women receive even fewer citations in emerging fields relative to their male counterparts: in emerging fields, female lead inventor patents receive an additional 1.4 to 1.8 fewer citations, suggesting that new fields exhibit a stronger pattern of undercitation for female inventors, rather than reducing barriers or bias.

4.3 Evolution of Undercitation Over the Sample Period

The estimates we present in the prior analyses suggest that across fields, patents with female lead inventors are consistently undercited relative to what would be expected for the same patent had its lead inventor been male. A natural question is whether these patterns vary over time, as gender norms, female participation in the workforce, and academia have evolved (Card, DellaVigna, Funk, and Iriberri (2022, 2023); Gompers and Wang (2017)).

Of course, for any given quality level, older patents may be more cited mechanically due to the increased passage of time allowing for citation. Without adjustments to our initial methodology, our findings may incorrectly suggest a decrease in undercitation over time, when in reality, it is simply a reflection of the fact that newer patents receive fewer citations on average simply by virtue of not having been around for as long. To ensure we are making appropriate comparisons, for the below analysis, we restrict the period during which forward citations are received to the first ten years post-patent grant. This method avoids the right censoring problem of forward citations, at the cost of excluding forward citations made after ten years out. While undercitation grows as the patent becomes more seasoned, as we show later in the robustness section, restricting to ten years allows us a clearer interpretation of results. Because we need to be able to measure ten years of forward citations, we necessarily must exclude from this analysis patents granted in the last decade of our sample. (We choose ten years in order to avoid excluding more than a decade of the sample period.) We thus create a sample of patents from 1976 to 2011 and measure the number of forward citations they receive within 10 years.

Plotting the gap in Figure 5 reveals no clear temporal trends. The gender-based citation gap initially increased in the first two decades but returned to baseline levels in the last decade. The effect's magnitude in the final year mirrors that of the initial years. These estimates do not necessarily suggest specific time trends or recent tendencies.

5 Who Undercites Female Inventors?

When applying for a patent, inventors cite supporting patents upon whose inventions the current patent builds. If the patent examiner determines that the inventor has omitted additional relevant citations, however, the examiner will add these to the patent application. Using this fact, we can better disentangle the source of undercitation: if it is driven by examiner-added citations or inventor-added citations, and whether the gender of the examiner or inventor plays a role.

To explore this, we first need to know which citations in a patent are attributable to the inventor versus the examiner. Starting in 2001, and more comprehensively starting in 2003, asterisks were added to the USPTO citation data to identify examiner-added patents in the data. Using this detail, we construct a new subsample starting from 2003, aggregating forward citations into four categories: (i) forward citations added (in a future patent) by male lead inventors, (ii) forward citations added by female lead inventors, (iii) forward citations added by male lead examiners, and (iv) forward citations added by female lead examiners. Using these groups, we can re-run our I-TEXT methodology and then decompose the sources of undercitation of female lead-inventor patents.

We begin our analysis by studying the gap resulting from the I-TEXT model for the four subgroups—citations that result from additions by male examiners, citations that result from additions by female examiners, citations that result from inclusion by male inventors, and citations that result from inclusion by female inventors. We present the gaps for the four subgroups in Figure 6. The gap for citations resulting from inclusion by male inventors is the largest among the subgroups, at -1.99. This is followed by the gap for citations resulting from inclusion by female lead inventors, with an estimate of -1.08. The gaps resulting from the inclusion of citations by examiners are smaller, with the gap for citations resulting from additions by female and male examiners of 0.019 and -0.2, suggesting more even-handedness from examiners than inventors.

We then turn to I-TEXT analysis, starting with examiner-added citations. We take all forward citations for a given patent that arise from the addition of the citation to a future patent application by an examiner. We then break these into forward citations added by female examiners and male examiners. Following a similar logic to our main tests, we then apply the I-TEXT model, estimating a neural net for male-lead inventor patents and a neural net for female-lead inventor patents. These networks allow us to estimate forward citation counts expected to be added by examiners of each gender, based on the gender of the lead inventor on the patent. We then calculate the $Delta$ between actual forward citations and the estimates of the examiner neural net for the same patent if it had a male lead inventor.

Table 4 presents the results of the estimation of regression models, using the I-TEXT derived $Delta$ as the dependent variable. Panel A presents estimates for female examiner-added forward citations, and Panel B presents the estimates for male examiner-added citations. The estimates in Panel A do not appear to indicate that female examiners, when adding citations to patents, *undercite* past female lead inventor patents. The coefficient estimates are economically small, ranging from 0.02 to 0.05 citations, with no statistical significance.

Panel B repeats the analysis for male examiners. Here, the estimates suggest that male examiners do undercite past female lead inventor patents in their additions to future patents. All estimated coefficients are statistically significant. The economic magnitudes are small, relative to the undercitation in our main results: the coefficient estimates range from -0.14 to -0.10 citations, with the economic magnitudes on the order of less than 1% of the unconditional sample mean. Overall, the estimates suggest a small undercitation of patents with a female lead inventor by male examiners in their additions to patents.

Having established this pattern, we next turn to citations added by future inventors to their patent applications. We conduct a similar analysis to that conducted above with examiners, focusing this time on the difference between the actual forward citations that stem from inclusion in future patents by male and female inventors and that which would be estimated for the same patents if their lead inventor had instead been male. The estimates from the I-TEXT *Delta* regressions are presented in Table 5. Panel A presents the results for female inventor-added forward citations, and Panel B presents the estimates for male inventor-added forward citations. The estimates across both panels suggest a clear pattern: both male and female inventors are significantly less likely to include a citation to a past female lead inventor patent than they would be estimated to have that same exact patent had a male lead inventor, but this gap is much larger for male lead inventors. The coefficient estimates for female inventor-added citations range from -0.38 to -0.52, with the strictest specifications statistically significant at conventional levels. For male inventor-added citations, the coefficient of interest ranges from -1.6 to -1.8 (approximately 8% of the unconditional sample mean), all statistically significant. Male inventors are significantly less likely to include a citation to a patent with a female lead inventor than they would be estimated to if that same patent had a male lead inventor.

Importantly, the quasi-random assignment of patent applications to examiners at the USPTO provides a useful setting that eases interpretation of these findings. Although not perfectly random, institutional assignment procedures yield substantial quasi-randomness within technology fields. Applications are initially allocated to technology-specific art units and subsequently assigned to examiners primarily based on workload and technology classification rather than applicant characteristics (Farre-Mensa, Hegde, and Ljungqvist (2020)). This assignment procedure resembles a "patent lottery," wherein examiner assignment is plausibly exogenous once routed within a technology group (Frakes and Wasserman (2017); Sampat and Williams (2019)).[18] While we cannot entirely eliminate selection on unobservables, controlling for art unit and examiner fixed effects substantially mitigates potential biases arising from systematic examiner characteristics, including, for example, variations in the representation of female examiners within art units handling more female lead inventor applications.

Given this, our estimates suggest that the undercitation of patents with a female lead inventor is driven in large part by male lead inventors under-citing existing female lead inventor patents in their own applications. It is important to note, however, that this undercitation does not necessarily imply discrimination on the part of male inventors. Alternative explanations

[18]Assignment of applications to examiners was randomized within art unit until 2019; see Aneja, Reshef, and Subramani (2024). Given this, we condition our sample prior to this date.

for this phenomenon could include men having more extensive familiarity with other male inventors, thus resulting in greater familiarity with patents filed by male lead inventors. Further research will be necessary to determine the underlying causes of this behavior and to distinguish between potential explanations such as discrimination or differences in professional networks and familiarity.

6 Robustness

6.1 Patent Significance

As implemented to this point, the I-TEXT approach inputs patent content together with patent novelty and patent impact. While breaking out the impact of any particular element of textual content from the neural network is challenging, we can partially open the black box to assess the impact of patent significance (impact/novelty). Specifically, we remove the novelty and impact inputs to the neural network and instead incorporate them in a single Significance measure into an OLS specification by estimating:

$$\begin{aligned} Delta_i &= \beta_1 I\left(FemaleInventor_i\right) + \beta_2 Significance_i + \beta_3 I\left(FemaleInventor_i\right) \times Significance_i \\ &\quad + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i, \end{aligned} \tag{10}$$

where $Delta_i$ is the patent-level citation premium, $I(FemaleInventor_i)$ is an indicator for a female lead inventor, and $Significance_i$ is the relative significance of patent i as defined in Equation 3. In this interacted OLS, β_1 identifies the gender gap in citations at $Significance = 0$; β_2 is the slope of citations with respect to significance for male-led patents (the baseline group); and β_3 captures how the gender gap varies with significance.

Estimates of Equation 10 are qualitatively similar to those in our baseline approach. Estimates in Table IA2 indicate that patents with female lead inventors are undercited by roughly 1.9 citations relative to what would have been estimated to have happened if the lead inventor on the same exact patent had instead been male. Notably, we also find that more significant patents are undercited more. Moreover, the interaction coefficient is negative; however, this coefficient loses statsitical significance when subjected to transforms to account for skewness in the distribution; see subsubsection 6.7.5.

6.2 Definition of "Female Inventor Patent"

A potential concern is that our results are attributable to how we classify patents by gender. In the analysis above, we use the name of the first inventor to assign gender to patents with multiple inventors. The first name on the patent is likely the most salient, as it is the first name observed when reading the patent. Of course, it is possible that examiners and inventors may consider all inventors or some subset of inventors when attributing authorship to the patent.

Importantly, our estimates are robust to several alternative approaches to attributing inventor gender to a patent. We follow the same estimation approach as in subsection 6.1. First, we limit our sample to patents with only one inventor, comparing female sole-inventor patents to male sole-inventor patents. This mitigates concerns that the gender of additional inventors other than the lead inventor may be driving the baseline results. The single-inventor sample also exhibits consistent undercitation of female inventor patents, as can be seen from Panel A of Table IA3. The estimates suggest that solo female inventors are undercited by nearly 2 patents per citation relative to what they would be expected to receive have they been male; nearly 9% of the sample mean. These results are highly significant and in line with the baseline results.

6.3 Alternative Embedding Models

The results presented up to this point utilize the patent texts and the Longformer embedding. However, a natural concern is that Longformer is not specific to (pre-trained on) scientific writing and, as a result, may not capture text quality for mediation purposes. To address this concern, we repeat our analyses using SciBERT embeddings. This embedding model does not make up our baseline approach because, like other BERT models, it has computational difficulties scaling to long texts. As a result, we utilize patent abstracts rather than longer patent texts for this exercise.

Table IA4 in the Internet Appendix presents the estimates using patent abstracts and specialized BERT embeddings. Estimates using SciBERT suggest that female lead inventor patents are undercited by 3.8 citations per patent relative to what they would have received had the lead inventor been male, which is equivalent to approximately 18% of the unconditional sample mean–a gap even larger than our baseline results. We note that these estimates result from neural networks that input the textual analysis of patent abstracts alone, and therefore utilize less information on patent content than the original Longformer-based analysis. That said, the results reinforce the patterns observed in the baseline results presented in Table 2.

6.4 Placebo Test

To further strengthen our conclusions, we run a placebo test. We adopt a randomized methodology, where a random sample of the dataset is selected and then segregated into two halves, one of which is randomly assigned to be "male" lead inventor and the other to be "female" lead inventor. We then apply the I-TEXT analytical procedure to ascertain the role, if any, that may be played by the methodology itself in shaping our results rather than the data. As expected, Table IA5 shows that, regardless of the model specification, there appears to be no statistical correlation between the random assignment to be a "female" patent and the resulting *Delta*. This is also reflected in Figure 7. All in all, the placebo tests suggest that the I-TEXT methodology itself is not likely to be driving our results.

6.5 Time Since Patent Grant

Does the undercitation of female lead inventor patents occur consistently over time from the patent grant date, or is it concentrated in the years immediately after the patent is granted, diminishing as inventors and examiners grow more familiar with a patent, or whether it intensifies over time, suggesting a self-reinforcing bias that becomes more challenging to address. For this analysis, we consider forward citations in distinct post-grant periods: [0-1) years, [1-5) years, [5-10) years, and [10-20] years. For each period, we gather forward citations to each patent in the sample and recalculate the *Delta*.

Table IA6 presents *Delta* estimates across periods. Over time, the *Delta* becomes more pronounced and statistically significant, indicating that undercitation may indeed be self-reinforcing. Specifically, the *Delta* estimates are -0.59, -0.60, and -2.2 for the [1-5) years, [5-10) years, and [10-20] years periods, respectively. These findings suggest that the gap of initial undercitation may exacerbate bias in later periods, potentially leading to a cycle where patents continue to be undercited.

6.6 Differences in Language by Gender

One consideration is that language may carry different meanings and be used differently by men and women. Prior research suggests that men describe their research's importance differently than women. If men emphasize genuine achievements more than women, this may give the impression that their patent is more significant, and lead to higher citations. Given this, we consider the potential impact of differences in language. First, as highlighted earlier, the I-TEXT framework helps address this issue directly by simultaneously analyzing two aspects: (1) it estimates the inventor's gender based on textual style, and (2) it employs a neural network to estimate citations based on the inventor's gender, allowing us to ensure that patents included in the analysis are not obviously identifiable as male- or female- lead inventor. We follow the approach of Veitch et al. (2020) and exclude patents where the estimated gender probability exceeds 97%, but our results are robust to the use of more restrictive cutoffs such as 95%, see Table IA7.

As a second approach, we re-estimate our regression models controlling for differences in tone and writing style of the patent. For each patent, we compute Number of Words, Sentiment Score (assessing the emotional tone of the text), Flesch-Kincaid and Flesch (readability tests indicating the difficulty of understanding the text), Gunning Fog (estimating the years of formal education needed to understand the writing), Coleman Liau (gauging the understandability of the text based on characters per word and sentence length), Dale Chall (assessing the comprehension difficulty based on sentence length and the percentage of difficult words), Ari (Automated Readability Index, similar to Flesch-Kincaid but using characters per word), Linsear Write (a readability formula developed for technical manuals, considering easy words, hard words, and long sentences), and Spache (a readability test for young readers, focusing on sentence length and unfamiliar words).[19] We then re-estimate our baseline specification while controlling for the additional language measures. We obtain qualitatively similar estimates, as can be seen in Table IA9. These estimates further reinforce that differences in writing styles do not appear to explain the differences in citation patterns between male and female lead inventors that we uncover.

[19]Summary statistics of these variables by gender can be found in Table IA8.

6.7 Additional Tests

6.7.1 Differences in Neural Network Error Rates

A consideration of the I-TEXT methodology is the possibility that differences in error rates between the male and female neural networks somehow drive the results. Delta is based on estimation errors for male-led patents and a combined error for female-led patents. In a scenario where the model for female-led patents is more accurate than that for male-led patents, the difference in the female case could turn negative, leading to a spurious result.

To rule out this concern, we compute the mean-squared errors for the two neural networks and conduct a Welch Two Sample t-test. If anything, the results suggest that the model for male-led patents is more accurate: we obtain a t-value of 1.496 (p-value of 0.1347), which points in the direction of the male model error being larger than the female model error. Overall, this suggests that the results are not being driven by the female model being more accurate than the male model.

6.7.2 Definition of Delta

While our main analysis considers the number of forward citations estimated under the male neural network as the baseline reference point, if our methodology is reasonable, we should be able instead to use the female neural network estimates as the baseline and observe an overcitation for male lead inventor patents in the data relative to if the inventor had instead been male. We confirm this by flipping our analysis, computing $Delta^{Female}$ as the difference between actual forward citations and the expected citations if the patent had a lead *female* inventor:

$$Delta_i^{Female} = ForwardCitation_i - Forward\widehat{Citation}_{i|T_i=1}, \tag{11}$$

where $ForwardCitation_i$ is the actual number of citations to a given patent with a lead inventor of a given gender and $Forward\widehat{Citation}_{i|T_i=1}$ is the number of citations implied by the I-TEXT model if the lead inventor had been female.

We then re-estimate our baseline model using $Delta^{Female}$ as our dependent variable and an indicator if the lead inventor is male. The estimates presented in Table IA10 align with the estimates using our original *Delta* measure, mitigating concerns that the specific definition of

Delta with the male reference point is somehow driving the undercitation bias measured in the paper.

6.7.3 Inventor Experience

Presumably, inventors with varying experience levels may receive different numbers of citations. Given this, we test whether the lead inventor's past experience, measured by the number of past patents, alters our baseline results. To do so, we include the number of past patents for the lead inventor in the set of inputs to our neural network model. We then use the resulting estimates to construct our measure of *Delta* and re-estimate our baseline model. Estimates in Table IA11 suggest that the undercitation documented in our main analysis persists even after accounting for the number of prior patents, suggesting our results are not driven by inventor experience.

6.7.4 Self-Citation

Another potential explanation for the observed gender bias in patent citations is that men and women may cite their own prior research at different rates. If there is a systematic difference in self-citation patterns based on gender, the results could be an artifact of this behavior rather than a true bias: if men cite their own patents more frequently, it could lead to the appearance of a citation bias. To address this, we start with our baseline sample and remove all patents that receive a forward *self-citation*. We then re-estimate our baseline specification using this updated sample. The results presented in Table IA12 suggest that our baseline findings are robust to the exclusion of self-citations. The estimates obtained using the updated sample are similar to those from our baseline analysis, strengthening our conclusion that the undercitation of patents with female lead inventors is not merely a result of differences in self-citation behavior between men and women.

6.7.5 Skewness in Citations

One potential consideration is that citations of patents exhibit skewness that our *Delta* measure may not account for. We address this by using an inverse hyperbolic sine transformation (asinh) for our *Delta* measure and re-estimating our baseline specification (*Delta* takes on both positive and negative values, limiting our ability to use a *log* transformation). The estimates presented in Table IA13 suggest an interpretation similar to our baseline estimates: patents with female lead inventors tend to be undercited relative to what would be expected had the lead inventor been male instead.[20]

6.7.6 Gender Propensity

In our baseline approach, we omit patents where the neural network had high confidence in the patent's gender. A concern, however, is that the removal of these patents may correlate with certain sub-fields that a gender is typically associated with Koning, Samila, and Ferguson (2021).

To address this, as a robustness test, we re-estimate our baseline specification using an expanded sample that includes patents for which the neural network had high confidence (greater than 97%) in identifying the gender of the lead inventor. Estimates in Table IA14 suggest that women remain similarly undercited in this expanded analysis. Both the economic magnitudes and statistical significance of the results are similar to our baseline estimates.

7 Economic Value of Patents and Citation Bias

If we assume that financial markets are efficient, the market value of a given patent should reflect its true underlying value: the present value of the expected free cash flows attributable to that patent. If our earlier findings of lower citation counts for female-led patents indeed stem from a systematic undercitation relative to their "true" quality, rather than an average lower quality of patent, then, controlling for observable characteristics, we should expect to see no economically meaningful difference in patent valuation by gender.

To explore this, we turn to Kogan et al. (2017)'s measure of patent value based on the stock market's response to key patent events (applications and grant announcements). Mechanically, this measure restricts our sample to patents assigned to publicly-traded firms. We then

[20] Our full set of reported results is robust to using the asinh transforation for *Delta*, available upon request. One notable difference when applying the asinh transformation is that the interaction coefficient for our tests of the effects of Significance, which is typically negative, becomes statistically indistinguishable from zero, as described in subsection 6.1 above. This attenuation of the (Significant×FemaleInventor) coefficient under the *asinh* specification is informative. It suggests that, in our main results, the stronger high-value gender gap observed in levels is driven primarily by additive tail effects rather than by proportional differences across the distribution. Notably, the main estimates for *FemaleInventor* remain both statistically and economically significant, underscoring the robustness of our findings.

explore whether the market sees different economic value for patents based on the gender of their lead inventors.

First, we replicate our base result–that patents with a female lead inventor are cited less than those with a male lead inventor–for this subsample. The estimates are presented in Table IA15. As in the full sample, we observe that female lead inventor patents are cited less than male lead inventor patents for patents that are assigned to publicly traded firms. The magnitude of the undercitation is similar to the full sample and represents an undercitation of approximately 12% relative to the sample mean.

Given this, we then turn to the market value of the patents. We estimate models with the following specification:

$$\begin{aligned} Dollar_i = \; & \beta_1 \cdot PatentWithFemaleLead_i \\ & + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i, \end{aligned} \tag{12}$$

where the variables and subscripts are defined as in previous sections. The dependent variable is the real dollar value (from Kogan et al. (2017), henceforth KPSS), and our main coefficient of interest, β_1, captures any differential in market valuation for patents led by female inventors relative to those led by male inventors. We double-cluster standard errors by examiner and patent-issue-year.

Panel A of Table 6 presents our baseline findings, including only the indicator for female lead inventor and patent characteristics fixed effects. Controlling for patent characteristics, the coefficient on *PatentWithFemaleLead* is economically small and statistically indistinguishable from zero. These estimates suggest that, controlling for patent attributes such as art unit, examiners, etc., the market assigns comparable value to patents with female lead inventors and those with male lead inventors, despite the lower citation counts for female lead inventor patents in the publicly-traded sample. These estimates are consistent with our interpretation that female lead inventor patents are undercited relative to their true quality (and what they would have received had the lead inventor been male instead.)

Next, we incorporated our citation measures. Specifically, we control both for log(Actual NumberofCitations) as well as log(ExpectedNumberofCitations $\mid T_i = 0$) – the number of citations the same patent would have been expected to receive based on our neural net model, had the lead inventor been male. Here, we estimate:

$$\begin{aligned} Dollar_i = \ & \beta_1 \cdot I\left(FemaleInventor_i\right) \\ & + \beta_2 \cdot log(ForwardCitations_i) \\ & + \beta_3 \cdot log(ExpectedNumberOfCitations_i | T_i = 0) \\ & + \beta_4 \cdot log(Novelty_i) \\ & + \beta_5 \cdot log(Impact_i) \\ & + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i. \end{aligned} \tag{13}$$

Panel B of Table 6 presents the estimates from the "horse race" model: across columns (1) through (5), the coefficient on $I\left(FemaleInventor_i\right)$ remains small and statistically insignificant. Furthermore, we observe no statistically significant loading in log(Forward Citations), while log(ExpectedNumberOfCitations $\mid T_i = 0$) loads positively and significantly throughout.[21] Additionally, the coefficient on novelty is positive and significant, while the coefficient on impact is statistically indistinguishable from zero. This is unsurprising, as novelty is known when the patent is granted, while impact is only determined 5 years after the patent is granted.

The estimates further reinforce our earlier interpretations: any observed differences in female- versus male-led patents' actual citation counts do not appear to translate into differential market valuations.

Collectively, these findings imply that citation counts may be a biased proxy for patent quality–potentially understating the contributions of female lead inventor patents. While our earlier results indicate that female lead inventors receive fewer citations than would be expected for the same exact patent content had the lead inventor instead been male, these lower counts do not appear to reflect lower economic value.

Importantly, the entity deriving "value" (as measured by the market values in KPSS) is the assignee corporation, and our estimates suggest that the value attributed to a patent for the assignee is not systematically different by the inventor's gender. Nevertheless, from the individual inventor's perspective, commonly accepted measures of patent quality–particularly citation counts–may still influence assessments of capability and, therefore, career outcomes. De-

[21] We additionally find that log(ExpectedNumberOfCitations $\mid T_i = 0$) also does well in explaining value in the upper tail of the distribution.

spite the lack of differences in *economic value*, female inventors might remain at a disadvantage–for example, when seeking promotions, opportunities at other firms, or investment in their startups–if the assessment of the quality of their patents is made based on citation counts alone.

8 Discussion and Conclusion

We provide evidence that patents with female lead inventors are undercited relative to what they would have received if their patent had a male lead inventor. Our approach uses new tools in machine learning to disentangle quality from forward citations, allowing us to show that the most commonly used measure for patent quality in fact under-recognizes the quality of patents with female inventors, relative to what that same patent would have received had the lead inventor been male.

Our findings have important economic implications. First, prior literature has highlighted that innovative activity is motivated by expected profits derived from the property rights granted to the patentee (Moser (2005, 2013)). If patents with female inventors are undercited, and as a result, female compensation for innovative labor is accordingly harmed, this may discourage women from entering the innovation economy. Such effects may further exacerbate the gender gap in STEM fields, Beede, Julian, Langdon, McKittrick, Khan, and Doms (2011).

A second important implication of our findings concerns the validity of research that relies on forward citations as a measure of patent quality. The existence of systematic gender-related biases in citations may lead to incorrect or misleading conclusions for research that relies on forward citations as a measure of patent quality. Given the large literature in economics, finance, and innovation that relies on forward citations as a proxy for quality, these findings suggest that a re-examination of relevant prior findings may be warranted.

Our paper also makes an important methodological contribution to the economic literature by introducing the I-TEXT methodology. Economics is steeped in the tradition of borrowing methodological innovations from adjacent fields. Big data, machine learning, and AI are new approaches that are poised to revolutionize empirical research in this field. Inference from Text Analysis can help researchers in answering key open economic questions. Our paper provides an initial roadmap for scholars to apply similar approaches in their own spheres.

References

Aghion, Philippe, John Van Reenen, and Luigi Zingales, 2013, Innovation and institutional ownership, *American Economic Review* 103, 277–304.

Aneja, Abhay, Oren Reshef, and Gauri Subramani, 2024, Attrition and the gender patenting gap, *Review of Economics and Statistics* 1–31.

Athey, Susan, and Guido W Imbens, 2019, Machine learning methods that economists should know about, *Annual Review of Economics* 11, 685–725.

Beede, David N, Tiffany A Julian, David Langdon, George McKittrick, Beethika Khan, and Mark E Doms, 2011, Women in stem: A gender gap to innovation, *Economics and Statistics Administration Issue Brief* .

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen, 2019, Who becomes an inventor in america? the importance of exposure to innovation, *Quarterly Journal of Economics* 134, 647–713.

Bellstam, Gustaf, Sanjai Bhagat, and J Anthony Cookson, 2021, A text-based analysis of corporate innovation, *Management Science* 67, 4004–4031.

Beltagy, Iz, Kyle Lo, and Arman Cohan, 2019, Scibert: A pretrained language model for scientific text, *arXiv preprint arXiv:1903.10676* .

Beltagy, Iz, Matthew E. Peters, and Arman Cohan, 2020, Longformer: The long-document transformer, *arXiv preprint arXiv:2004.05150* .

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri, 2020, Are referees and editors in economics gender neutral?, *Quarterly Journal of Economics* 135, 269–327.

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri, 2022, Gender differences in peer recognition by economists, *Econometrica* 90, 1937–1971.

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri, 2023, Gender gaps at the academies, *Proceedings of the National Academy of Sciences* 120, e2212421120.

Chawla, Dalmeet Singh, 2016, Men cite themselves more than women do, *Nature* 535, 212.

Comenetz, Joshua, 2016, Frequently occurring surnames in the 2010 census, *United States Census Bureau* 1–8.

Cong, Lin William, Tengyuan Liang, and Xiao Zhang, 2019, Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information, *Interpretable, and Data-driven Approach to Analyzing Unstructured Information (September 1, 2019)* .

Desai, Pranav, 2019, Biased regulators: Evidence from patent examiners, Working paper.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, 2018, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* .

Erel, Isil, Léa H Stern, Chenhao Tan, and Michael S Weisbach, 2021, Selecting directors using machine learning, *Review of Financial Studies* 34, 3226–3264.

Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist, 2020, What is a patent worth? evidence from the us patent “lottery”, *Journal of Finance* 75, 639–682.

Farre-Mensa, Joan, Zack Liu, and Jordan Nickerson, 2022, Do startup patent acquisitions affect inventor productivity?, Working paper.

Frakes, Michael D, and Melissa F Wasserman, 2017, Is the time allocated to review patent applications inducing examiners to grant invalid patents? evidence from microlevel application data, *Review of Economics and Statistics* 99, 550–563.

Gavrilova, Evelina, and Steffen Juranek, 2021, Female inventors: The drivers of the gender patenting gap, Working paper.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019a, Text as data, *Journal of Economic Literature* 57, 535–74.

Gentzkow, Matthew, Jesse M Shapiro, and Matt Taddy, 2019b, Measuring group differences in high-dimensional choices: method and application to congressional speech, *Econometrica* 87, 1307–1340.

Gompers, Paul A, and Sophie Q Wang, 2017, Diversity in innovation, Technical report, National Bureau of Economic Research.

Graham, Stuart JH, Alan C Marco, and Richard Miller, 2018, The uspto patent examination research dataset: A window on patent processing, *Journal of Economics & Management Strategy* 27, 554–578.

Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg, 2005, Market value and patent citations, *RAND Journal of Economics* 16–38.

Hanley, Kathleen Weiss, and Gerard Hoberg, 2019, Dynamic interpretation of emerging risks in the financial sector, *Review of Financial Studies* 32, 4543–4603.

Hansen, Stephen, Michael McMahon, and Andrea Prat, 2018, Transparency and deliberation within the fomc: a computational linguistics approach, *Quarterly Journal of Economics* 133, 801–870.

Hengel, Erin, and Euyoung Moon, 2020, Gender and equality at top economics journals, Working paper.

Hirschey, Mark, and Vernon J Richardson, 2004, Are scientific indicators of patent quality useful to investors?, *Journal of Empirical Finance* 11, 91–107.

Hirshleifer, David, Po-Hsuan Hsu, and Dongmei Li, 2013, Innovative efficiency and stock returns, *Journal of Financial Economics* 107, 632–654.

Hunt, Jennifer, 2016, Why do women leave science and engineering?, *ILR Review* 69, 199–226.

Iaria, Alessandro, Carlo Schwarz, and Fabian Waldinger, 2024, Gender gaps in academia: Global evidence over the twentieth century, *Available at SSRN 4150221* .

Jensen, Kyle, Balázs Kovács, and Olav Sorenson, 2018, Gender differences in obtaining and maintaining patent rights, *Nature Biotechnology* 36, 307–309.

Jha, Manish, Hongyi Liu, and Asaf Manela, 2025, Does finance benefit society? a language embedding approach, *The Review of Financial Studies* hhaf012.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2021, Measuring technological innovation over the long run, *American Economic Review: Insights* 3, 303–20.

Khandelwal, Anant, 2021, Fine-tune longformer for jointly predicting rumor stance and veracity, in *Proceedings of the 3rd ACM India Joint International Conference on Data Science & Management of Data (8th ACM IKDD CODS & 26th COMAD)*, 10–19.

Khetan, Vivek, Roshni Ramnani, Mayuresh Anand, Subhashis Sengupta, and Andrew E Fano, 2022, Causal bert: Language models for causality detection between events expressed in text, in *Intelligent Computing*, 965–980 (Springer).

Koffi, Marlène, 2021, Gendered citations at top economic journals, in *AEA Papers and Proceedings*, volume 111, 60–64, American Economic Association.

Koffi, Marlène, 2024, Innovative ideas and gender inequality, Working paper.

Koffi, Marlène, and Matt Marx, 2023, Cassatts in the attic, Technical report, National Bureau of Economic Research.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665–712.

Koning, Rembrand, Sampsa Samila, and John-Paul Ferguson, 2021, Who do we invent for? patents by women focus more on women's health, but few women get to invent, *Science* 372, 1345–1348.

Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan, 2021, Measuring corporate culture using machine learning, *Review of Financial Studies* 34, 3265–3315.

Li, Yikuan, Ramsey M Wehbe, Faraz S Ahmad, Hanyin Wang, and Yuan Luo, 2022, Clinical-longformer and clinical-bigbird: Transformers for long clinical sequences, *arXiv preprint arXiv:2201.11838* .

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov, 2019, Roberta: A robustly optimized bert pretraining approach, *arXiv preprint arXiv:1907.11692* .

Loughran, Tim, and Bill McDonald, 2016, Textual analysis in accounting and finance: A survey, *Journal of Accounting Research* 54, 1187–1230.

Martinez, Gema Lax, Julio Raffo, Kaori Saito, et al., 2016, *Identifying the gender of PCT inventors*, volume 33 (WIPO).

Moser, Petra, 2005, How do patent laws influence innovation? evidence from nineteenth-century world's fairs, *American Economic Review* 95, 1214–1236.

Moser, Petra, 2013, Patents and innovation: evidence from economic history, *Journal of Economic Perspectives* 27, 23–44.

National Science Foundation, 2013, Women, minorities, and persons with disabilities in science and engineering: 2013, Accessed: 2024-09-23.

Oster, Emily, 2019, Unobservable selection and coefficient stability: Theory and evidence, *Journal of Business & Economic Statistics* 37, 187–204.

Reshef, Oren, Abhay Aneja, and Gauri Subramani, 2021, Persistence and the gender innovation gap: evidence from the us patent and trademark office, in *Academy of Management Proceedings*, volume 2021, 11626, Academy of Management Briarcliff Manor, NY 10510.

Rouen, Ethan, Kundal Sachdeva, and Aaron Yoon, 2022, The evolution of esg reports and the role of voluntary standards, *Available at SSRN* 4227934.

Routledge, Bryan R, Stefano Sacchetto, and Noah A Smith, 2017, Predicting merger targets and acquirers from text, Working paper.

Sampat, Bhaven, and Heidi L Williams, 2019, How do patents affect follow-on innovation? evidence from the human genome, *American Economic Review* 109, 203–36.

Sarsons, Heather, 2017, Recognition for group work: Gender differences in academia, *American Economic Review* 107, 141–45.

Sarsons, Heather, Klarita Gërxhani, Ernesto Reuben, and Arthur Schram, 2021, Gender differences in recognition for group work, *Journal of Political Economy* 129, 101–147.

Shao, Yifan, Haoru Li, Jinghang Gu, Longhua Qian, and Guodong Zhou, 2021, Extraction of causal relations based on sbel and bert model, *Database* 2021.

Sherman, Mila Getmansky, and Heather E Tookes, 2022, Female representation in the academic finance profession, *Journal of Finance* 77, 317–365.

Toole, A. A., M. Saksena, C. A. W. deGrazia, K. Black, F. Lissoni, E. Miguelez, and G. Tarasconi, 2020, Progress and potential: 2020 update, Technical report, U.S. Patent and Trademark Office.

Tzioumis, Konstantinos, 2018, Demographic aspects of first names, *Scientific Data* 5, 1–9.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, 2017, Attention is all you need, *Advances in Neural Information Processing Systems* 30.

Veitch, Victor, Dhanya Sridhar, and David Blei, 2020, Adapting text embeddings for causal inference, in *Conference on Uncertainty in Artificial Intelligence*, 919–928, PMLR.

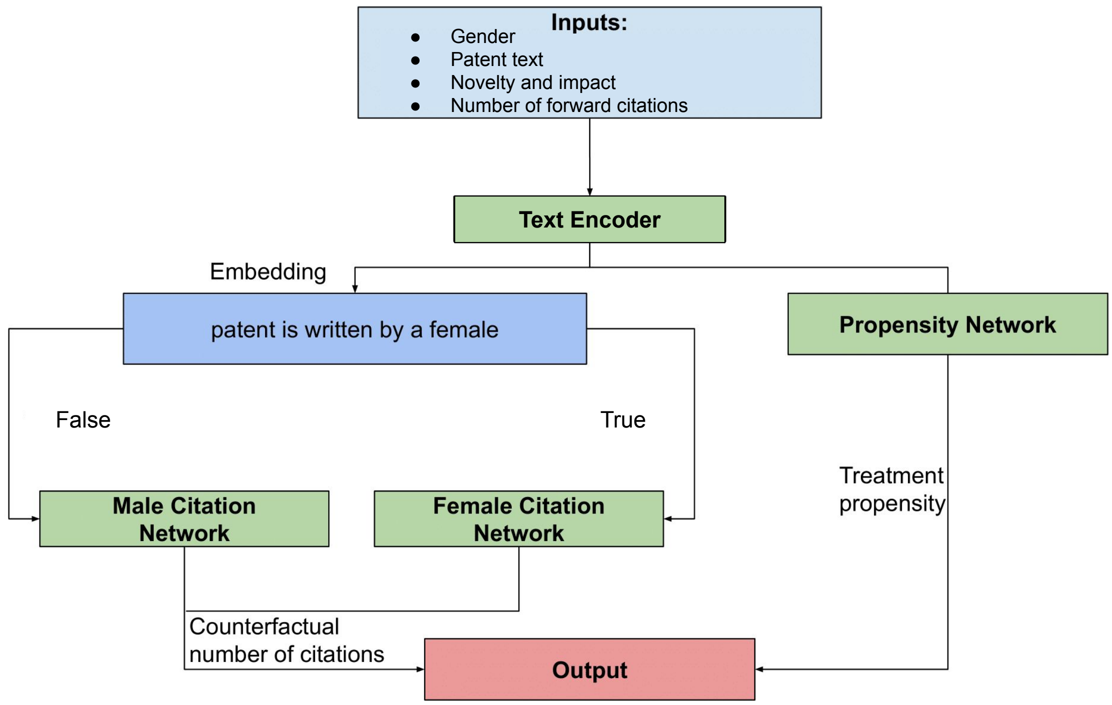

FIGURE 1: I-TEXT ESTIMATION PROCEDURE

The figure illustrates the estimation procedure of I-TEXT once the neural networks are trained. The light blue block at the top describes the input used for estimation. The green blocks are the four neural networks trained using the patent data. The blue block describes the decision rule used for outcome estimation. Finally, the red block is the output that combines the outputs of the citation estimation networks and the propensity score estimation network.

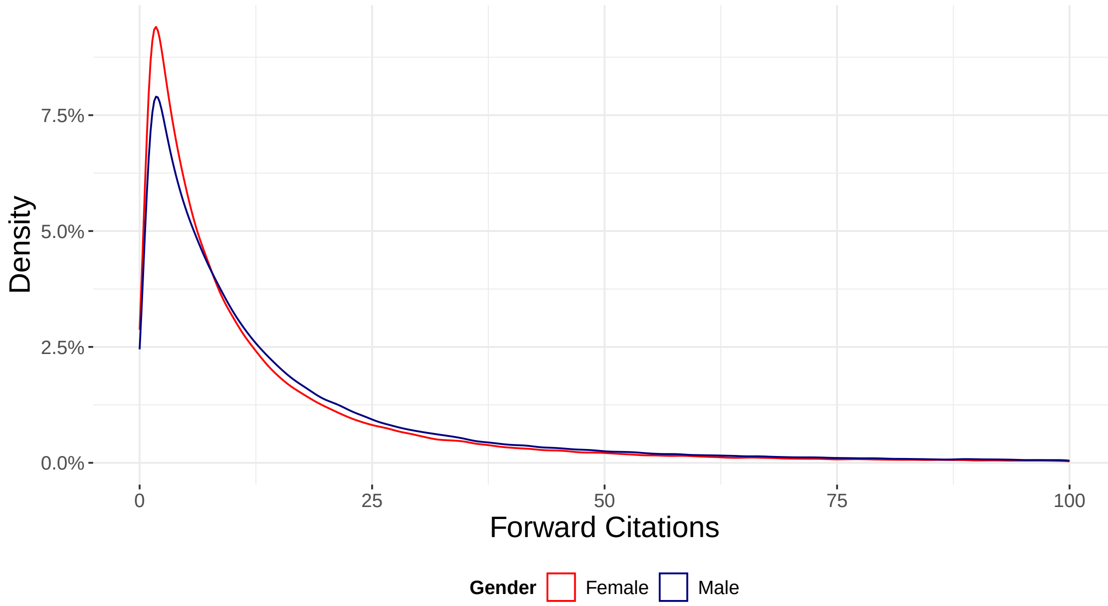

(a) Panel A: Observed Forward Citations

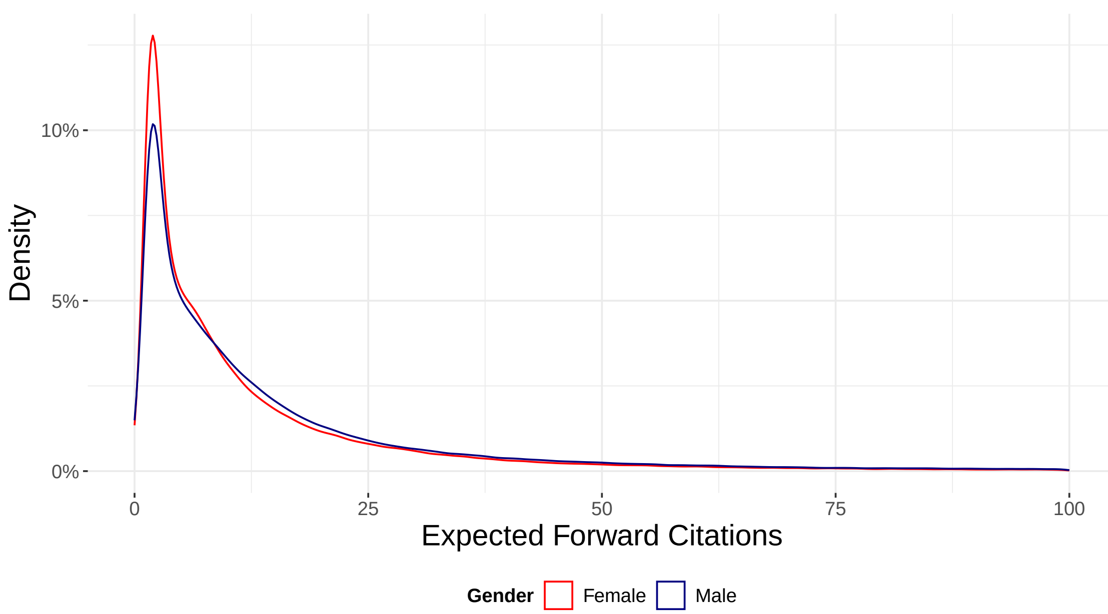

(b) Panel B: I-TEXT Implied Forward Citations

FIGURE 2: DISTRIBUTION OF FORWARD CITATIONS

This figure illustrates the distribution of forward citations. Panel A uses forward citations observed in the data, while Panel B uses the expected number of forward citations as implied by I-TEXT. The horizontal axis counts the number of citations, while the vertical axis measures the percent of the distribution. The red line corresponds to females. The blue line corresponds to males. The distribution is truncated at 100 for ease of interpretation. The natural logarithm transformation of these distributions is presented in Figure IA2.

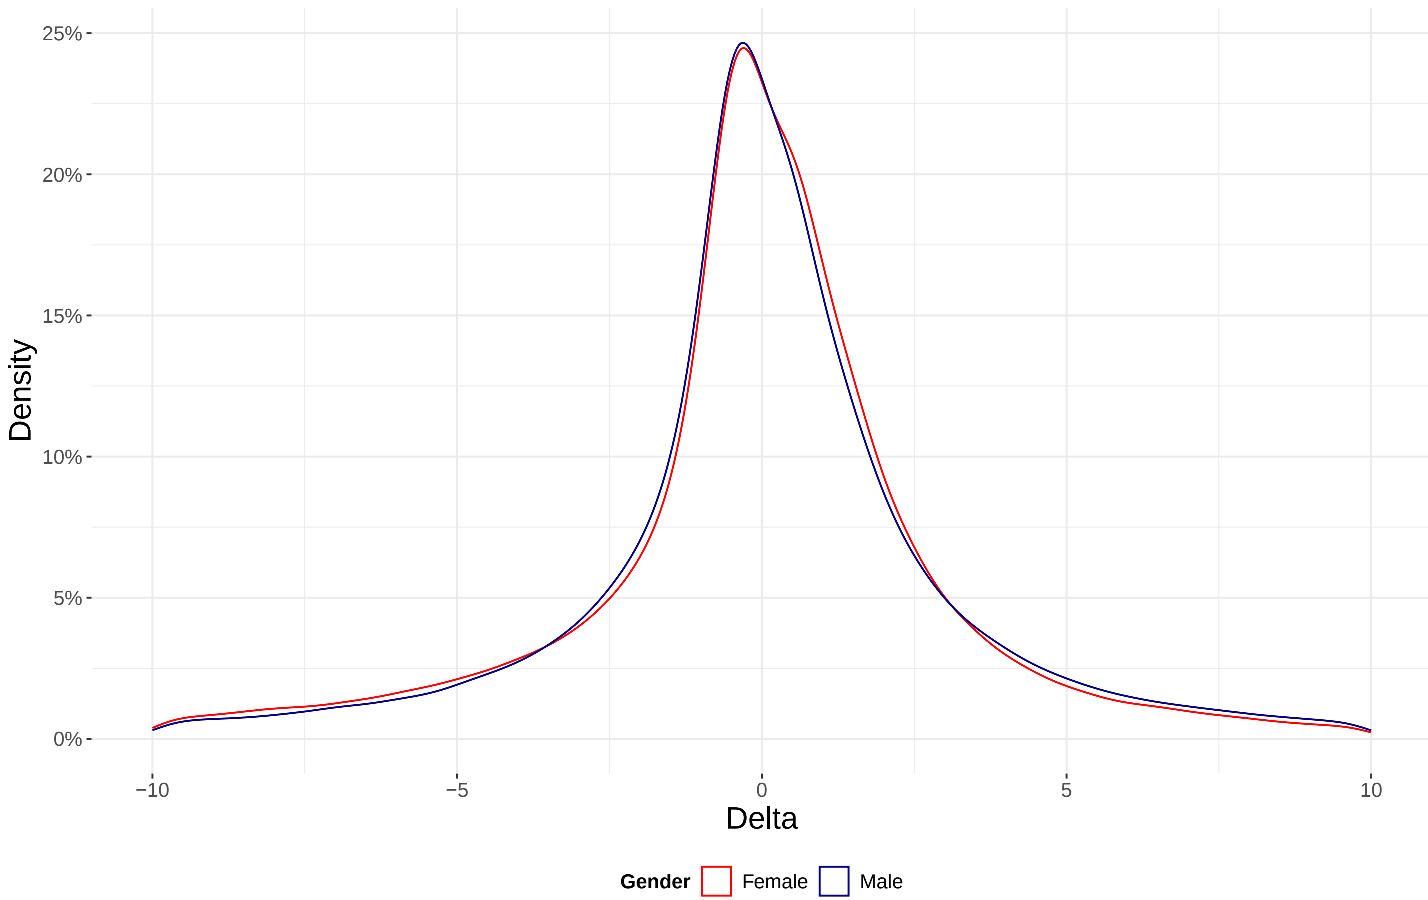

FIGURE 3: DISTRIBUTION OF DELTA

This figure illustrates the difference between forward citations and expected forward citations for male lead inventor patents, as defined by Equation 8. The red line corresponds to females, the blue line corresponds to males. The distribution is truncated between -10 to 10.

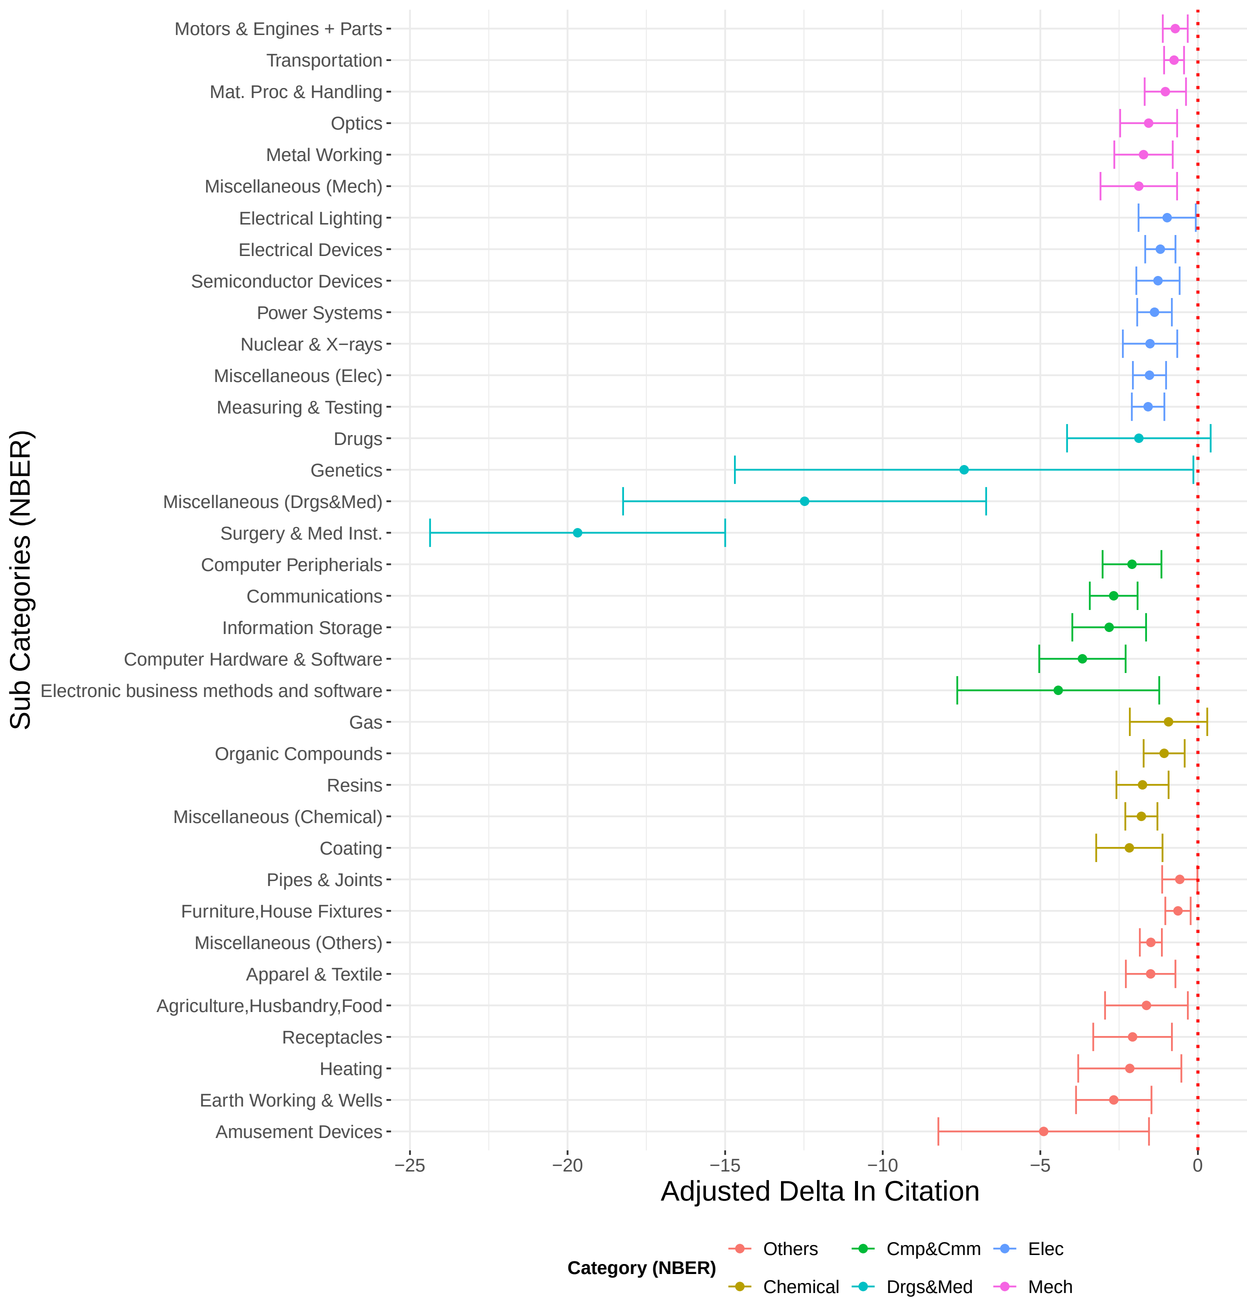

FIGURE 4: DELTA IN CITATIONS BY PATENT SUB CATEGORIES

This figure illustrates the coefficients of Equation 9. For ease of interpretation, each point corresponds to the linear combination of the baseline result for females and the interaction terms. Whiskers correspond to a 95% confidence interval. Coefficients are sorted by patent category and then by the magnitude of the estimate. Colors correspond to the patent category as defined by the NBER, where blue observations correspond to mechanical (Mech), light blue corresponds to electrical (Elec), green observations correspond to drugs and medical (Drgs&Med), yellow observations correspond to computers and communication (Cmp&Comm), red observations correspond to chemical (Chemical), pink observations correspond to other (Other) categories. The red dotted line is plotted at the zero intercept, representing no effect.

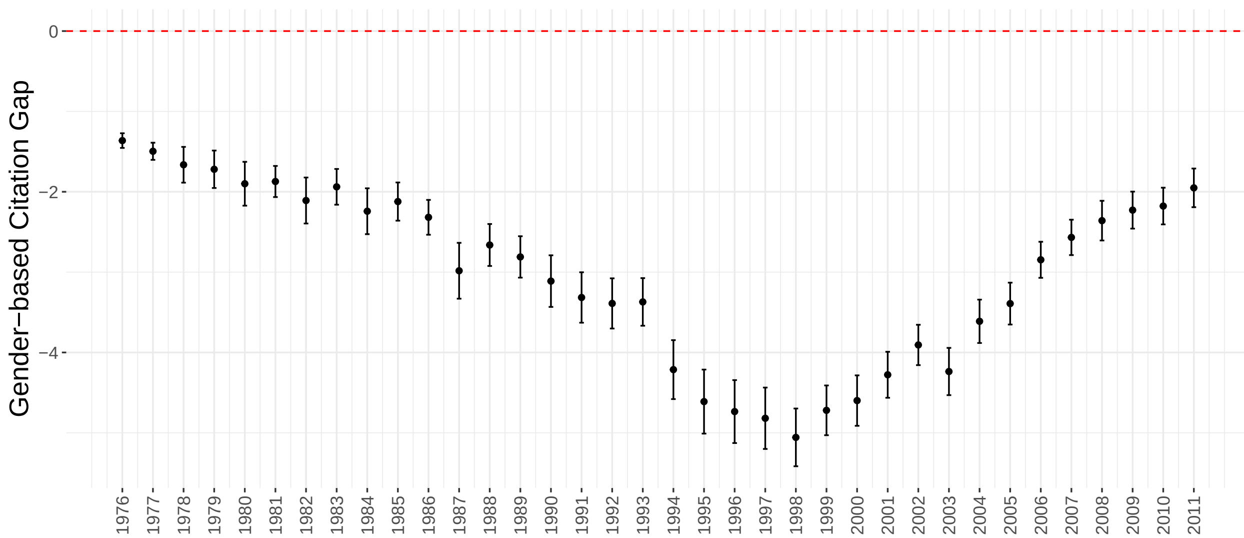

FIGURE 5: EVOLUTION OF GENDER-BASED CITATION GAP OVER TIME

This figure illustrates the evolution of citations over time. The horizontal axis corresponds to the year a patent was granted. The vertical axis corresponds to the delta in forward citations, with negative numbers corresponding to undercitation. Forward citations are computed within the first ten years the patent was granted and uses the I-TEXT methodology. Each point represents an estimate from a separate estimate, with error bars corresponding to a 95% confidence interval. All estimates include customer, examiner, and examiner art unit fixed effects.

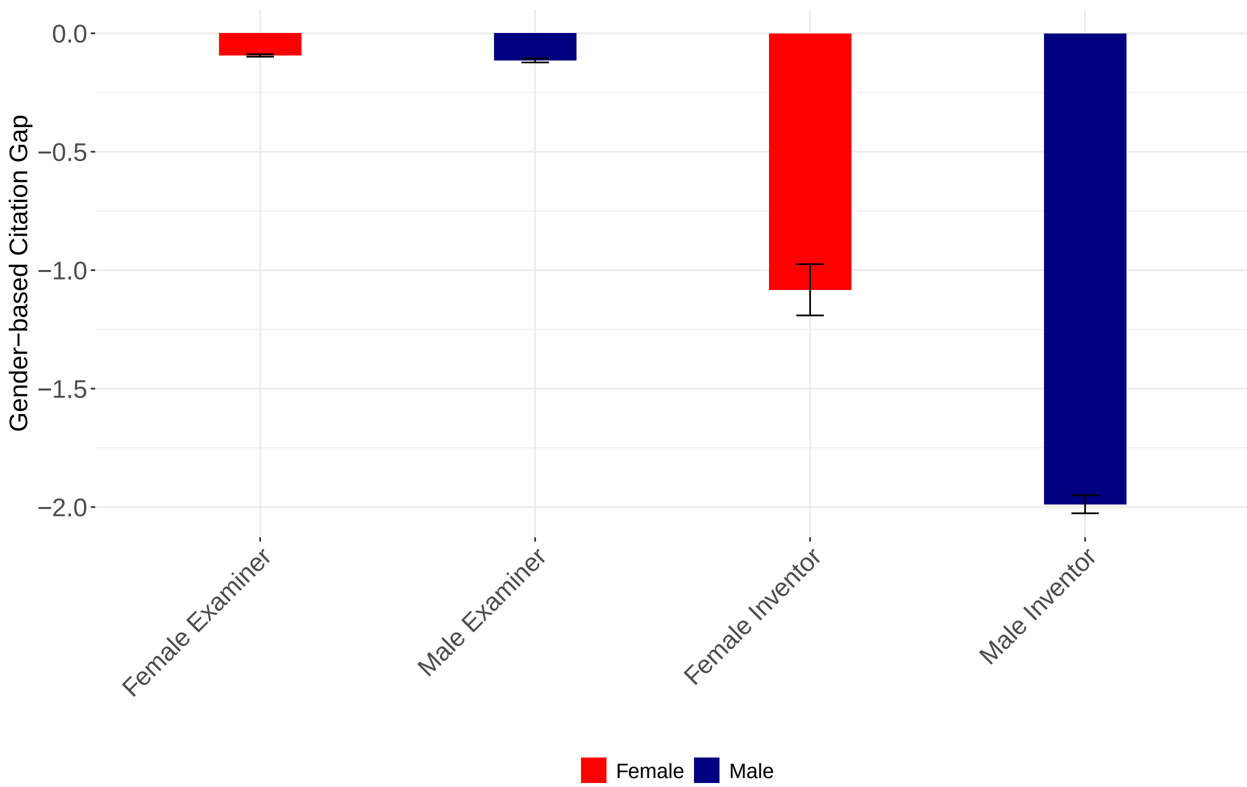

FIGURE 6: GENDER-BASED CITATION GAP, EXAMINER AND INVENTOR ADDED CITATIONS

This figure illustrates the delta by examiner and inventor added citations and by gender. The red columns correspond to female-added citations, and the blue columns correspond to male-added citations. Error bars correspond to the 95% confidence interval.

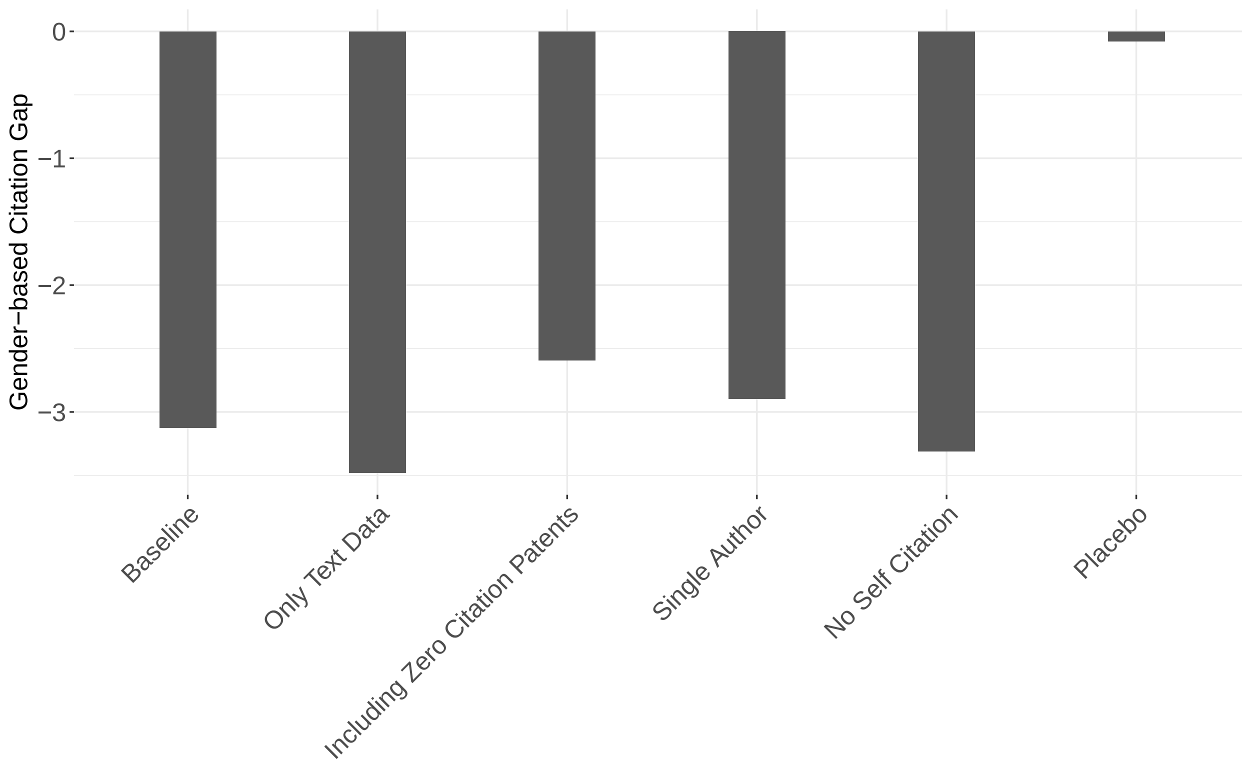

FIGURE 7: GENDER-BASED CITATION GAP FOR DIFFERENT SAMPLES

This figure plots the gender-based citation gap as defined by Equation 6. Each bar represents a different sample and estimate. The first bar corresponds our baselines specification. The second bar corresponds to the I-TEXT methodology using only text data. The third bar expands the sample to includes patents with zero citations. The fourth bar corresponds to patents with a single author. The fourth bar corresponds to teams with a majority of a given gender. The fifth bar excludes patents that self-cite. The sixth bar is a placebo test.

TABLE 1: SUMMARY STATISTICS

This table provides summary statistics on patents and citations. The sample covers patents issued from 1976 through 2015. Panel A presents a two-way table of forward citations by gender. Panel B presents a two-way table of patents in the top decile by gender. Panel C presents a two-way table of patents by their cooperative patent classification (CPC). ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data Source: USPTO.

Gender of Lead Inventor	Male			Female			
	N	Mean	SD	N	Mean	SD	Test
Panel A: Difference in Forward Citations							
Forward Citation	289,847	23.62	62.39	232,120	19.2	55.17	F=715.64***
Significance (standardized)	289,847	0.04	0.97	232,120	-0.05	1.03	F=933.537***
Panel B: Difference By Top Decile Innovations							
Breakthrough	289,847			232,120			χ^2=985.594***
→ No	258,237	89%		212,828	92%		
→ Yes	31,610	11%		19,292	8%		
Panel C: Difference by Cooperative Patent Classification							
CPC Section	289,847			232,120			χ^2=3860.740***
→ Chemistry	29,895	10%		28,058	12%		
→ Electricity	64,772	22%		60,268	26%		
→ Fixed Constructions	9,967	3%		5,544	2%		
→ Human Necessities	36,276	13%		29,996	13%		
→ Mechanical Engineering	27,683	10%		15,752	7%		
→ Performing Operations	52,116	18%		33,770	15%		
→ Physics	65,775	23%		56,473	24%		
→ Textiles	3,333	1%		2,242	1%		

TABLE 2: LEAD FEMALE INVENTORS

This table estimates the relationship between the gender of the lead inventor and citations. The dependent variable is the expected forward citation for each patent. This specification updates the I-TEXT architecture to include the significance measure with the text embeddings. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the patent art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: No Adjustment

	Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−3.294***	−2.985***	−2.877***	−2.44***	−1.51***
	(0.411)	(0.396)	(0.381)	(0.33)	(0.22)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	521 967	521 967	521 967	521 967	521 967
R^2	0.096	0.106	0.166	0.265	0.484

Panel B: Text Adjustment

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.251***	−2.281***	−2.331***	−2.462***	−2.468***
	(0.295)	(0.296)	(0.309)	(0.312)	(0.291)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	521 967	521 967	521 967	521 967	521 967
R^2	0.033	0.034	0.079	0.170	0.350

TABLE 3: EMERGING FIELDS

This table studies the citations to new fields of innovation. *Emerging Field* takes the value of one if the art unit first appeared within five years of the patent being granted. Panel A uses the forward citation for each patent as its dependent variable. Panel B uses I-Text adjustment, with the dependent variable being the difference in the observed number and the expected number of citations for a patent if the lead inventor was male, as defined by Equation 8. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: No Adjustment

	Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Emerging Field	11.24***	5.01***	3.92***	3.82***	2.444**
	(2.49)	(1.67)	(1.17)	(1.13)	(0.995)
Lead Female Inventor	−3.382***	−3.064***	−2.957***	−2.482***	−1.583***
	(0.407)	(0.401)	(0.378)	(0.334)	(0.254)
Emerging Field × Lead Female Inventor	−0.128	−0.802	−0.810	−1.301	0.218
	(0.890)	(0.854)	(0.783)	(0.997)	(0.903)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	494 239	494 239	494 239	494 239	494 239
R^2	0.098	0.107	0.167	0.267	0.488

Panel B: Text Adjustment

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Emerging Field	−0.675	0.0112	0.627	0.717	0.778*
	(0.488)	(0.4777)	(0.501)	(0.465)	(0.430)
Lead Female Inventor	−2.221***	−2.250***	−2.308***	−2.436***	−2.45***
	(0.294)	(0.295)	(0.308)	(0.305)	(0.29)
Emerging Field × Lead Female Inventor	−1.755**	−1.687**	−1.588**	−1.713**	−1.448**
	(0.732)	(0.731)	(0.748)	(0.741)	(0.593)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	494 239	494 239	494 239	494 239	494 239
R^2	0.033	0.034	0.079	0.171	0.352

TABLE 4: EXAMINER GENDER COUNTERFACTUALS

This table studies the source of examiner-added citations for male inventors. The dependent variable is the difference in forward citations. The dependent variable is *Delta*, the difference in the observed number and the expected number of citations for a patent if the lead author was male, as defined by Equation 8. Panel A uses the difference in forward citations that were added by female lead examiners as its dependent variable. Panel B uses the difference in forward citations that were added by male lead examiners as its dependent variable. The sample covers patents issued from 1976 through 2015. Note, the source of citations is only available following the start of 2001. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Citation Added by Female Lead Examiner

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.1185***	−0.1101***	−0.1149***	−0.1374**	−0.0689
	(0.0246)	(0.0249)	(0.0393)	(0.0523)	(0.1398)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	17 603	17 603	17 603	17 603	17 603
R^2	0.089	0.094	0.414	0.671	0.799

Panel B: Citation Added by Male Lead Examiner

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.1111***	−0.1049***	−0.1028***	−0.0842**	−0.0679
	(0.0276)	(0.0281)	(0.0307)	(0.0331)	(0.0415)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	72 505	72 505	72 505	72 505	72 505
R^2	0.030	0.033	0.204	0.400	0.561

TABLE 5: INVENTOR-ADDED CITATIONS

This table studies the source of inventor-added citations. The dependent variable is the difference in forward citations. The dependent variable is $Delta$, the difference in the observed number and the expected number of citations for a patent if the lead inventor was male, as defined by Equation 8. Panel A uses the difference in forward citations that were added by female lead inventors as its dependent variable. Panel B uses the difference in forward citations that were added by male lead inventors as its dependent variable. The sample covers patents issued from 1976 through 2015. Note, the source of citations is only available following the start of 2001. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Citation Added by Female Lead Inventors

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.384	−0.479*	−0.496	−0.301 701***	−0.515 619***
	(0.279)	(0.250)	(0.739)	(0.000 001)	(0.000 001)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	4727	4727	4727	4727	4727
R^2	0.279	0.294	0.761	0.923	0.954

Panel B: Citation Added by Male Lead Inventors

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−1.627***	−1.629***	−1.649***	−1.760***	−1.696***
	(0.238)	(0.238)	(0.247)	(0.237)	(0.257)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	266 015	266 015	266 015	266 015	266 015
R^2	0.021	0.022	0.078	0.214	0.355

TABLE 6: VALUE OF PATENT

This table presents the relationship between the gender of the lead inventor and the real dollar value of the granted patent. The dependent variable is deflated to 1982 (million) dollars using the CPI, as calculated in Kogan et al. (2017). The sample covers patents with a positive number of citations, granted from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Without Controls

	Dollar Value				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.41**	−0.26*	−0.27*	−0.15	−0.12
	(0.17)	(0.15)	(0.14)	(0.11)	(0.15)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	191 086	191 086	191 086	191 086	191 086
R^2	0.090	0.115	0.188	0.450	0.593

Panel B: With Citation and Novelty Controls

	Dollar Value				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.21	−0.17	−0.19	−0.12	−0.12
	(0.15)	(0.14)	(0.14)	(0.11)	(0.15)
log(Actual Citations)	0.17	0.061	0.095	0.029	−0.11
	(0.22)	(0.173)	(0.178)	(0.171)	(0.11)
log(Expected Male)	1.84***	1.55***	1.46***	0.80***	0.38**
	(0.46)	(0.37)	(0.39)	(0.29)	(0.16)
log(novelty)	2.62***	1.01**	1.28**	0.92**	0.39*
	(0.86)	(0.42)	(0.51)	(0.41)	(0.21)
log(impact)	1.10	−0.27	−0.0062	0.24	0.29
	(0.74)	(0.27)	(0.3482)	(0.32)	(0.22)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	191 075	191 075	191 075	191 075	191 075
R^2	0.098	0.119	0.191	0.451	0.593

INTERNET APPENDIX
FOR ONLINE PUBLICATION

This internet appendix provides supplemental materials for the paper.

- Section A contains the additional figures and tables used to supplement the main results of our paper.
- Section B explains the details of BERT starting from a simple neural network.
- Section C explains the encoder architecture and our I-TEXT model in terms of how inferences are made from texts.
- Section E discusses the potential concern of the first-mover advantage and how it makes our estimates conservative estimates.
- Section F shows word clouds that display words that are correlated with a patent being estimated as having low or high text quality.
- Section G discusses the derivation of *Delta* from an econometrics standpoint.
- Section H discusses the construction of different subsets used in different tests and the sample size for each test.

Appendix A Additional Figures and Tables

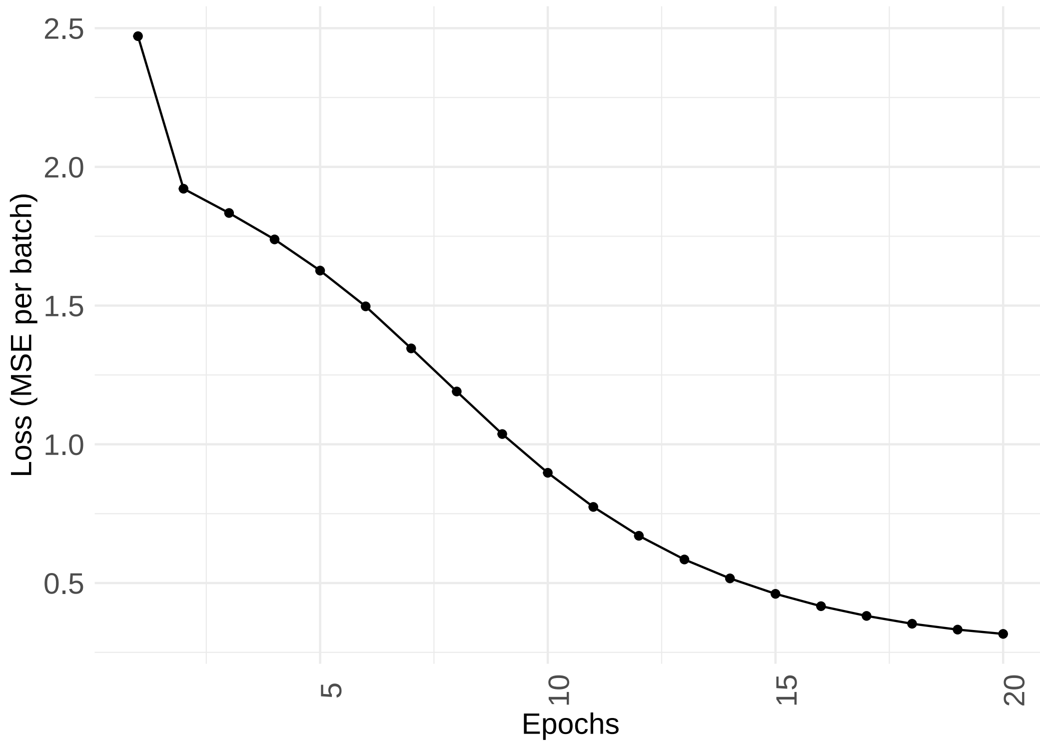

FIGURE IA1: LOSS FUNCTION

This figure illustrates the loss function of the I-TEXT model. The horizontal axis corresponds to the number of complete passes of the training dataset through the algorithm or epoch. The vertical axis corresponds to the loss function and is the mean square error per batch.

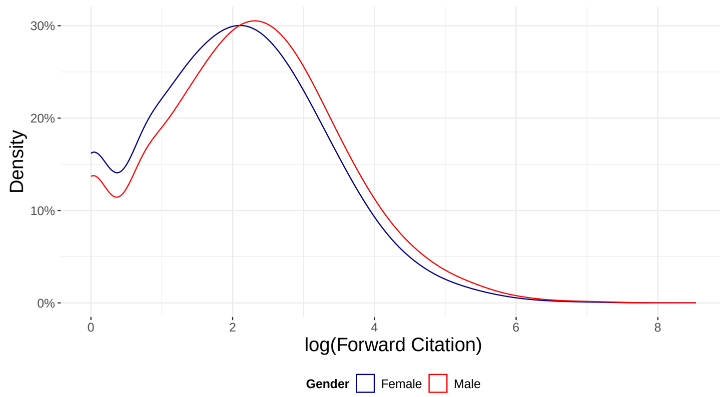

(a) Forward Citations

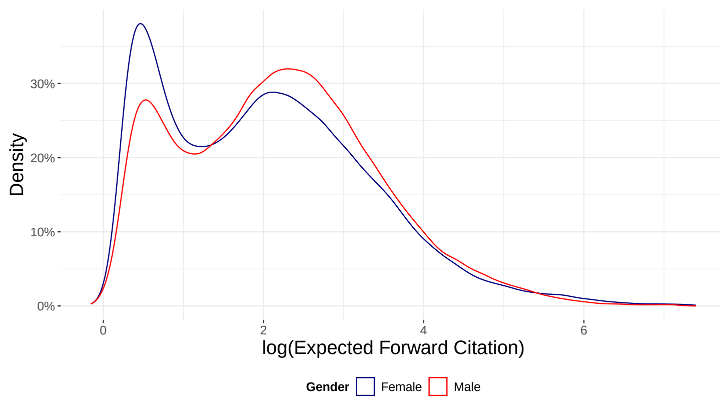

(b) Model Implied Forward

FIGURE IA2: DISTRIBUTION OF FORWARD CITATIONS

This figure illustrates the transformation from forward citations to expected forward citations. Panel A uses the natural logarithm of forward citations while Panel B uses the natural logarithm of forward citations expected from our model. The vertical axis in both panels measures the percent of the distribution. The red line corresponds to females, and the blue corresponds to males.

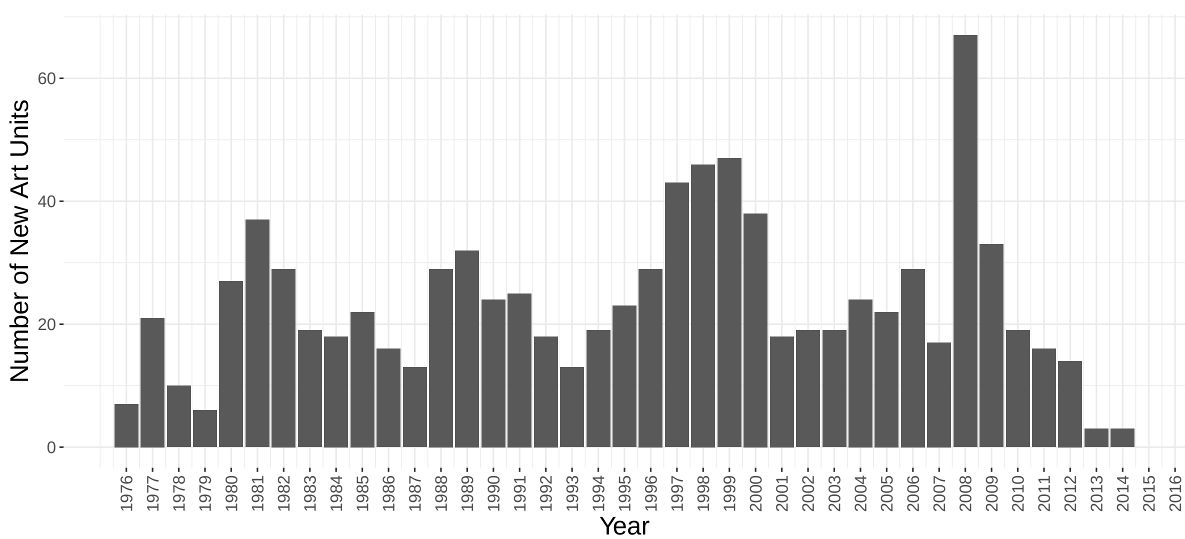

FIGURE IA3: NEW FIELD OVER TIME

This figure illustrates the number of emerging fields that arise over time. Emerging fields are measured by the first time the art unit appears for an examiner.

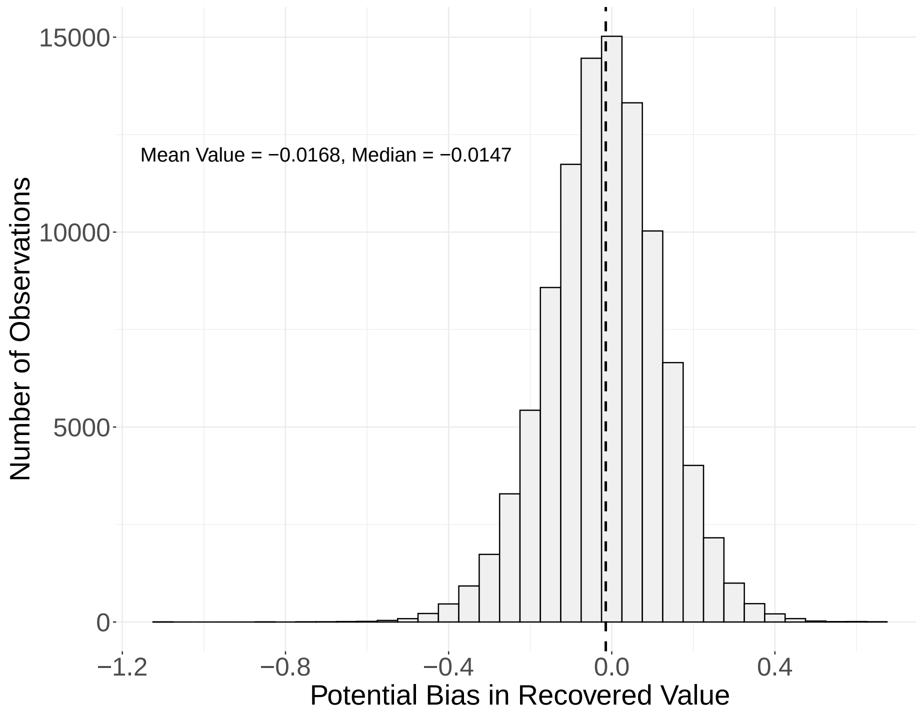

FIGURE IA4: DISTRIBUTION OF ESTIMATED MALE AND FEMALE CITATIONS IN THE SYNTHETIC TEST

The figure plots the distribution of errors from the synthetic tests. The zero axis indicates no bias from the recovered value. The mean error is -0.0168, while the median value is -0.0147.

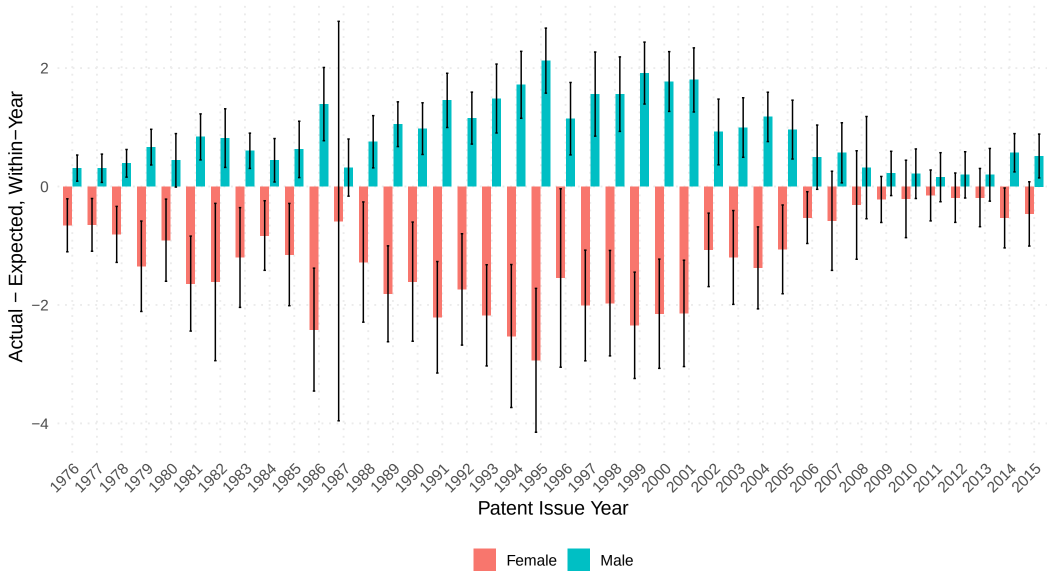

Panel A: Delta over time

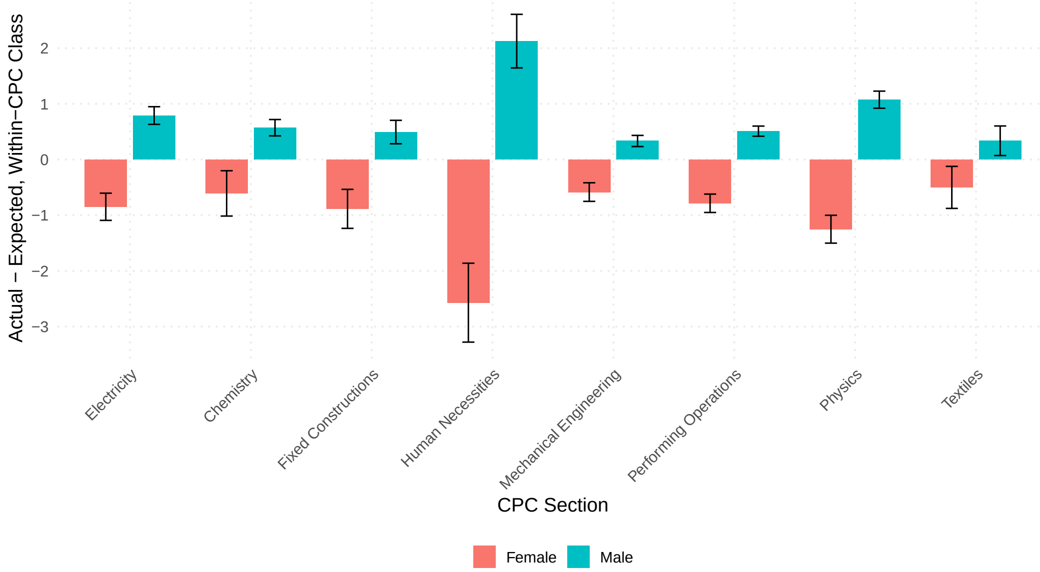

Panel B: Citation distribution by CPC category

FIGURE IA5: EXPOSITION OF DELTA MEASURE

This figure compares citation outcomes of female-led and male-led patents. Panel A reports the average citation *Delta* by gender over time, while Panel B reports the average citation *Delta* across CPC technology classes. In both panels, red bars denote female-led patents and blue bars denote male-led patents. Error bars indicate 95% confidence intervals for the difference between actual and expected citations.

TABLE IA1: POPULATION CONSIDERATIONS

This table estimates the relationship between the gender of the lead inventor and the delta in forward citations. For both panels, the training set is conducted on a balanced subset. Panel A estimates *delta* using the full sample. Panel B uses the balanced subset and weights observations by the underlying data generating process. The table estimates cover all patents granted from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Full Sample

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.10*	−2.25**	−2.27**	−2.40**	−2.02*
	(1.08)	(1.03)	(1.03)	(1.01)	(1.01)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	3 889 731	3 889 731	3 889 731	3 889 731	3 889 731
R^2	0.019	0.020	0.034	0.084	0.212

Panel B: Weighted Regression

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.14***	−2.160***	−2.174***	−2.277***	−2.267***
	(0.29)	(0.289)	(0.299)	(0.307)	(0.304)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	521 967	521 967	521 967	521 967	521 967
R^2	0.027	0.028	0.077	0.169	0.358

TABLE IA2: UNDERCITATION OF LEAD FEMALE INVENTOR, I-TEXT

This table estimates the relationship between the gender of the lead inventor and citations using the I-Text methodology. Panel A uses a balanced subset of patents with a positive number of citations. Panel B uses a balanced subset of all patents. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the patent art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Intensive Margin

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−1.740***	−1.725***	−1.746***	−1.882***	−2.009***
	(0.139)	(0.139)	(0.135)	(0.124)	(0.127)
Significance	−2.135***	−2.78***	−2.635***	−2.480***	−1.960***
	(0.285)	(0.44)	(0.456)	(0.466)	(0.406)
Lead Female Inventor × Significance	−1.074***	−1.081***	−1.184***	−1.173***	−1.071***
	(0.205)	(0.221)	(0.223)	(0.219)	(0.217)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	554 486	554 486	554 486	554 486	554 486
R^2	0.045	0.045	0.087	0.167	0.348

Panel B: Extensive Margin

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−1.896***	−1.891***	−1.905***	−1.990***	−1.969***
	(0.144)	(0.145)	(0.145)	(0.145)	(0.134)
Significance	−1.346***	−1.609***	−1.494***	−1.34***	−0.948***
	(0.173)	(0.306)	(0.307)	(0.29)	(0.249)
Lead Female Inventor×Significance	−1.134***	−1.130***	−1.188***	−1.143***	−1.06***
	(0.194)	(0.193)	(0.194)	(0.208)	(0.19)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	738 386	738 386	738 386	738 386	738 386
R^2	0.023	0.023	0.056	0.136	0.274

TABLE IA3: ROBUSTNESS TO GENDER DEFINITION

This table studies the sensitive of results to the nature of the inventor team. The dependent variable is *Delta*, the difference in the observed number and the expected number of citations for a patent if the lead inventor were male, as defined by Equation 8. This table uses a balanced subset of single-inventor patents. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.309***	−2.30***	−2.336***	−2.351***	−1.991***
	(0.224)	(0.23)	(0.247)	(0.231)	(0.226)
Significance	−0.56***	−1.425***	−1.113***	−1.28***	−1.35***
	(0.17)	(0.358)	(0.284)	(0.33)	(0.33)
Lead Female Inventor × Significance	−2.190***	−2.283***	−2.449***	−2.374***	−2.049***
	(0.477)	(0.506)	(0.517)	(0.495)	(0.429)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	193 725	193 725	193 725	193 725	193 725
R^2	0.028	0.029	0.157	0.291	0.529

TABLE IA4: ROBUSTNESS TO LANGUAGE MODEL

This table replaces the Longformer model with an alternative language model (SciBERT). The sample covers patents issued from 1976 through 2015. The dependent variable is *Delta*, the difference in the observed number and the expected number of citations for a patent if the lead inventor was male. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−3.907***	−3.90***	−3.903***	−3.934***	−3.806***
	(0.283)	(0.28)	(0.284)	(0.276)	(0.263)
Significance	−2.464***	−3.203***	−2.943***	−2.616***	−2.090***
	(0.363)	(0.573)	(0.576)	(0.529)	(0.491)
Lead Female Inventor × Significance	−2.637***	−2.652***	−2.725***	−2.667***	−2.442***
	(0.398)	(0.396)	(0.402)	(0.416)	(0.368)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	654 552	654 552	654 552	654 552	654 552
R^2	0.053	0.054	0.096	0.186	0.346

TABLE IA5: PLACEBO TEST

This table presents a placebo test by randomizing the gender of patents and re-running our I-TEXT approach to establish that the association are not an artifact of I-TEXT. The dependent variable is *Delta*, the difference in the observed number and the expected number of citations for a patent if the lead inventor were male, as defined by Equation 8. The sample covers patents with a positive number of citations, granted from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.117	−0.115	−0.054	−0.0684	−0.0222
	(0.146)	(0.146)	(0.139)	(0.1375)	(0.1646)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	204 695	204 695	204 695	204 695	204 695
R^2	0.010	0.010	0.089	0.226	0.435

TABLE IA6: YEARS AFTER PATENT IS GRANTED

This table estimates Equation 7 and studies the difference in forward citations by the number of years after the patent was granted. The dependent variable is *Delta*, the difference in the observed number and the expected number of citations for a patent if the lead inventor was male, as defined by Equation 8. Column (1) – (4), study the difference in forward citations 0-1, 2-5, 6-10, and 11-20 years after they are granted, respectively. All specifications use *Art Unit*, *Patent Grant Year*, *Examiner*, *Attorney*, and *Assignee* fixed effects. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level.***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta			
	0-1 Years	2-5 Years	6-10 Years	11-20 Years
	(1)	(2)	(3)	(4)
Lead Female Inventor	0.073	−0.588***	−0.597***	−2.18***
	(0.591)	(0.049)	(0.091)	(0.29)
Significance	−0.039	0.078*	0.11	0.077
	(0.886)	(0.041)	(0.14)	(0.143)
Lead Female Inventor × Significance	0.012	−0.083	−0.026	−1.68***
	(0.326)	(0.076)	(0.076)	(0.26)
Art Group FE	Yes	Yes	Yes	Yes
Patent Issue Year FE	Yes	Yes	Yes	Yes
Examiner FE	Yes	Yes	Yes	Yes
Attorney FE	Yes	Yes	Yes	Yes
Assignee FE	Yes	Yes	Yes	Yes
Observations	11 637	244 104	259 436	308 207
R^{2}	0.787	0.339	0.230	0.416

TABLE IA7: ROBUSTNESS TO CUTOFFS

This table tests the robustness of our baseline result by using 5% and 95% cutoff. The dependent variable is *Delta*, the difference in the observed number and the expected number of citations for a patent if the lead inventor were male, as defined by Equation 8. The sample covers patents with a positive number of citations, granted from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.057***	−2.091***	−2.148***	−2.267***	−2.261***
	(0.273)	(0.274)	(0.287)	(0.286)	(0.265)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	472 332	472 332	472 332	472 332	472 332
R^2	0.034	0.035	0.084	0.184	0.376

TABLE IA8: DIFFERENCE IN WRITING STYLES

This table reports the difference in the writing style between males and females. The sample covers patents issued from 1976 through 2015 and has at least 120 words. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO, Google Patents.

Gender	Male			Female			
	N	Mean	SD	N	Mean	SD	Test
Flesch Kincaid	316631	13.21	2.87	236942	13.42	3.09	F= 633.614***
Flesch	316631	42.64	11.55	236942	41.7	12.09	F= 863.66***
Gunning Fog	316631	16.71	3.06	236942	16.87	3.29	F= 337.912***
Coleman Liau	316631	11.23	1.97	236942	11.39	2.03	F= 844.521***
Dale Chall	316631	11.58	1.3	236942	11.71	1.38	F= 1453.23***
Ari	316631	13.61	3.67	236942	13.88	3.96	F= 684.596***
Linsear Write	316631	16.95	4.69	236942	17.21	5.18	F= 361.431***
Spache	316631	8.61	1.11	236942	8.7	1.21	F= 764.935***
Sentiment	316631	0.87	0.33	236942	0.87	0.33	F= 0.067
Sentiment Score	316631	0.05	0.06	236942	0.05	0.06	F= 70.679***
Word Count	316631	4028.97	5388.38	236942	4022.51	5634.83	F= 0.187

TABLE IA9: DELTA IN FORWARD CITATION, LANGUAGE CONTROLS

This table estimates the relationship between the gender of the lead inventor and citations. The dependent variable is the *Delta* in forward citations. This specification includes controls for the Number of Words, Sentiment Score, Flesch-Kincaid, Flesch, Gunning Fog, Coleman Liau, Dale Chall, Ari, Linsear Write, and Spache measure of the patent's text. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the patent art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−1.92***	−1.93***	−1.94***	−2.00***	−2.09***
	(0.16)	(0.16)	(0.15)	(0.14)	(0.13)
Significant	−2.35***	−2.6***	−2.48***	−2.36***	−1.89***
	(0.27)	(0.4)	(0.42)	(0.43)	(0.39)
Significant × Lead Female Inventor	−1.05***	−1.04***	−1.14***	−1.17***	−1.08***
	(0.21)	(0.22)	(0.22)	(0.22)	(0.22)
Language Controls	Yes	Yes	Yes	Yes	Yes
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	553 569	553 569	553 569	553 569	553 569
R^2	0.048	0.048	0.090	0.169	0.349

TABLE IA10: ALTERNATIVE DEFINITION OF DELTA

This table considers $Delta^{Female}$ as the dependent variable. The dependent variable is the difference between the actual number of forward citations and the expected number of citations if the patent had a lead female inventor. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Significance Included in Neural Network

	$Delta^{Female}$				
	(1)	(2)	(3)	(4)	(5)
Lead Male Inventor	2.434***	2.44***	2.451***	2.440***	2.196***
	(0.265)	(0.27)	(0.275)	(0.276)	(0.264)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	521 967	521 967	521 967	521 967	521 967
R^2	0.007	0.007	0.048	0.132	0.302

Panel B: Controlling for Significance

	$Delta^{Female}$				
	(1)	(2)	(3)	(4)	(5)
Lead Male Inventor	1.435***	1.439***	1.443***	1.450***	1.330***
	(0.103)	(0.104)	(0.102)	(0.102)	(0.107)
Significance	−1.514***	−1.527***	−1.606***	−1.605***	−1.35***
	(0.249)	(0.307)	(0.306)	(0.316)	(0.31)
Lead Male Inventor × Significance	1.085***	1.095***	1.15***	1.181***	1.148***
	(0.144)	(0.144)	(0.15)	(0.145)	(0.152)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	554 486	554 486	554 486	554 486	554 486
R^2	0.012	0.013	0.052	0.122	0.292

TABLE IA11: PATENT EXPERIENCE

This table estimates the relationship between the gender of the lead inventor and citations. The dependent variable is the expected forward citation for each patent. This specification updates the I-TEXT architecture to include the lead inventor's prior patent count with the text embeddings. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the patent art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.922***	−2.947***	−2.993***	−3.155***	−3.188***
	(0.382)	(0.382)	(0.402)	(0.391)	(0.375)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	550 198	550 198	550 198	550 198	550 198
R^2	0.049	0.050	0.095	0.187	0.378

TABLE IA12: OMITTING SELF-CITATION

This table considers the role of self-citations. The dependent variable is the expected forward citation for each patent. The sample omits patents that self cite and covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: No Adjustment

	Forward Citation				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−1.584***	−1.592***	−1.564***	−1.163***	−0.540**
	(0.331)	(0.319)	(0.317)	(0.293)	(0.213)
Significance	10.96***	12.1***	11.46***	10.71***	9.06***
	(1.09)	(1.5)	(1.47)	(1.42)	(1.19)
Lead Female Inventor × Significance	−1.494**	−1.445**	−1.297**	−1.371**	−1.210*
	(0.576)	(0.626)	(0.609)	(0.623)	(0.622)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	462 717	462 717	462 717	462 717	462 717
R^2	0.115	0.117	0.175	0.272	0.485

Panel B: Text Adjustment

	Delta in Forward Citation				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−1.70***	−1.699***	−1.728***	−1.831***	−1.869***
	(0.13)	(0.131)	(0.129)	(0.126)	(0.127)
Significance	−2.272***	−2.955***	−2.840***	−2.671***	−2.123***
	(0.283)	(0.439)	(0.454)	(0.468)	(0.446)
Lead Female Inventor × Significance	−0.880***	−0.883***	−1.010***	−1.0***	−0.969***
	(0.176)	(0.200)	(0.199)	(0.2)	(0.217)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	462 717	462 717	462 717	462 717	462 717
R^2	0.045	0.046	0.094	0.182	0.373

TABLE IA13: DELTA IN FORWARD CITATION, ASINH TRANSFORMATION

This table transforms the outcome variable *Delta* using the inverse hyperbolic functions (asinh). The sample focuses on patents with a positive number of forward citations. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	asinh(Delta in Forward Citations)				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.1600***	−0.1596***	−0.162***	−0.1753***	−0.186***
	(0.0128)	(0.0127)	(0.013)	(0.0128)	(0.014)
Significance	−0.1586***	−0.1801***	−0.1649***	−0.1553***	−0.1264***
	(0.0179)	(0.0249)	(0.0258)	(0.0252)	(0.0246)
Lead Female Inventor × Significance	−0.0066	−0.006 41	−0.009 57	−0.006 34	−0.004 92
	(0.0115)	(0.012 11)	(0.011 84)	(0.012 17)	(0.013 32)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	554 486	554 486	554 486	554 486	554 486
R^2	0.033	0.033	0.066	0.140	0.280

TABLE IA14: INCLUDING PATENTS WITH HIGH GENDER PROPENSITY

This table estimates the relationship between the gender of the lead inventor and the delta in forward citations. Sample does not exclude patents for which the probability of being associated with a certain gender is high. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.599***	−2.629***	−2.642***	−2.726***	−2.724***
	(0.258)	(0.261)	(0.267)	(0.268)	(0.262)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	687 830	687 830	687 830	687 830	687 830
R^2	0.028	0.029	0.065	0.141	0.296

TABLE IA15: SUBSAMPLE TEST, PUBLICALLY TRADED FIRMS

This table estimates the relationship between the gender of the lead inventor and citations using the I-Text methodology. The sample conditions on patents where the assignees are publically traded firms. For Panel A, the dependent variable is the number of forward citations. For Panel B, the dependent variable is *Delta*, the difference in the observed number and the expected number of citations for a patent if the lead inventor were male, as defined by Equation 8. The sample covers patents with a positive number of citations, granted from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.76***	−2.42***	−2.44***	−2.22***	−1.57***
	(0.44)	(0.44)	(0.46)	(0.49)	(0.38)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	191 086	191 086	191 086	191 086	191 086
R^2	0.133	0.143	0.243	0.380	0.508

	Delta in Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−2.91***	−2.94***	−2.94***	−2.98***	−2.92***
	(0.35)	(0.35)	(0.35)	(0.35)	(0.32)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	191 086	191 086	191 086	191 086	191 086
R^2	0.039	0.040	0.119	0.245	0.337

TABLE IA16: ALTERNATIVE SPECIFICATION

This table consider an alternative specification, using the expected number of citations as a control variable. The dependent variable is number of forward citations. The estimate uses a Poisson specification with high-dimensional fixed effects. The sample covers patents issued from 1976 through 2015. Standard errors are clustered at the art unit and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Forward Citations				
	(1)	(2)	(3)	(4)	(5)
Lead Female Inventor	−0.1752***	−0.1668***	−0.1649***	−0.1393***	−0.116 98***
	(0.0173)	(0.0168)	(0.0159)	(0.0115)	(0.008 77)
Art Group FE	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	Yes
Observations	521 967	521 967	521 967	521 967	521 967
Adj. Pseudo R2	0.491	0.518	0.558	0.620	0.700

Appendix B Explanation of neural networks and BERT

One of the major tools we use are text transformers which convert texts to numerical vectors. Here we present an introduction to BERT (the base model for the set of models we test). We start from a basic definition of a one-layer neural network – the first-level building blocks of these models. Then, we discuss how multiple layers are grouped together to produce contextually meaningful text embedding vectors through the self-attention mechanism. Lastly, we discuss how the BERT class of models are built and trained.

B.1 Single-layer neural networks

We start by discussing a simple neural network: a one-hidden-layer linear neural network. This network is defined by three dimensions: input dimension $\dim_{in}$, hidden dimension $\dim_h$, and output dimension $\dim_{out}$.

There are two mappings in this one-hidden-layer network. The first is mapping from the input space to the hidden space. Mathematically, let the input data be X

$$H = f_1(X) = XW_1 + B_1$$

where W_1 and B_1 are trainable matrices of parameters.

The second function maps from the hidden space to the output space

$$Y = f_2(H) = f_2(f_1(X)) = (XW_1 + B_1)W_2 + B_2$$

$$= XW_1W_2 + B_1W_2 + B_2$$

where W_2 and B_2 are trainable matrices of parameters.

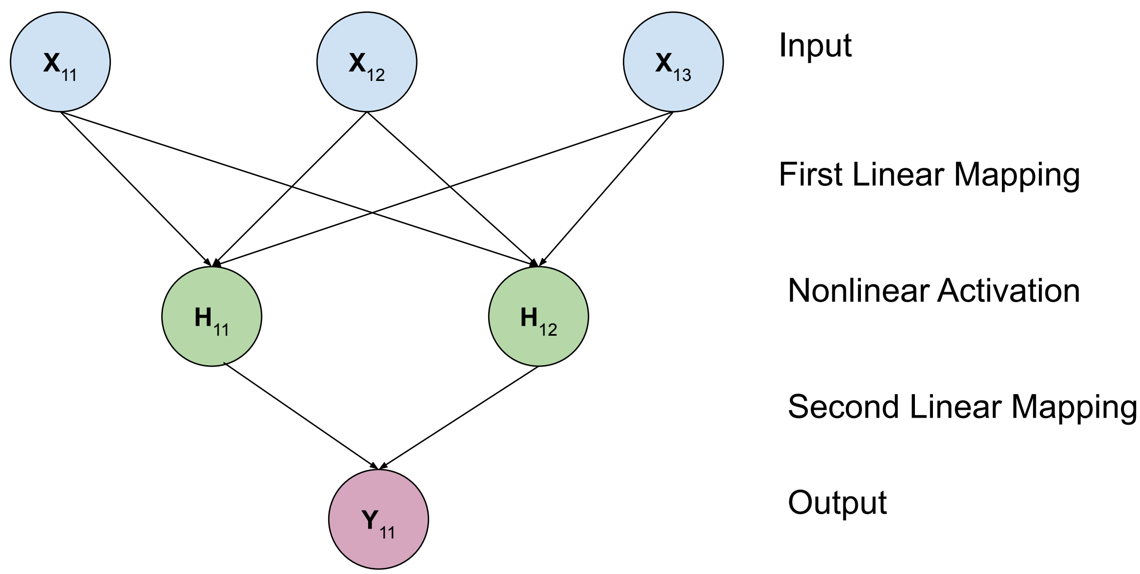

FIGURE IB1: DIAGRAM OF A ONE-HIDDEN LAYER NEURAL NETWORK WITH A NONLINEAR ACTIVATION FUNCTION

As shown in Figure IB1, this simple linear network can be generalized by adding a nonlinear activation function $g(\cdot)$. When we apply the activation function to the first mapping, the output of the first mapping becomes

$$H = g(f_1(X)) = g(XW_1 + B_1)$$

Similarly, we can apply a nonlinear activation function after any mapping to add nonlinearity to a linear transformation, allowing neural networks to more flexibly fit any relation between the input X and output Y.

B.2 Self-attention

The goal of self-attention is to create a numerical embedding for each piece of text, and this embedding respects each token's contextual relation with all of the tokens in the text. More specifically, the raw input to the attention mechanism is a piece of text T. Then, this text is broken into sub-word tokens in a parsing process called tokenization. This set of tokens is pre-defined such that a relatively limited number of tokens can be combined to represent a large number of unique words. For example, the prefix "un" is a token in many models because it has a meaning of negation when combined with many other sub-word tokens like "happy". Many other tokens are short and common words like "and".

After tokenization, each token is assigned a naive embedding that combines a representation of the meaning of the word and the position of the word in the whole text. The result is

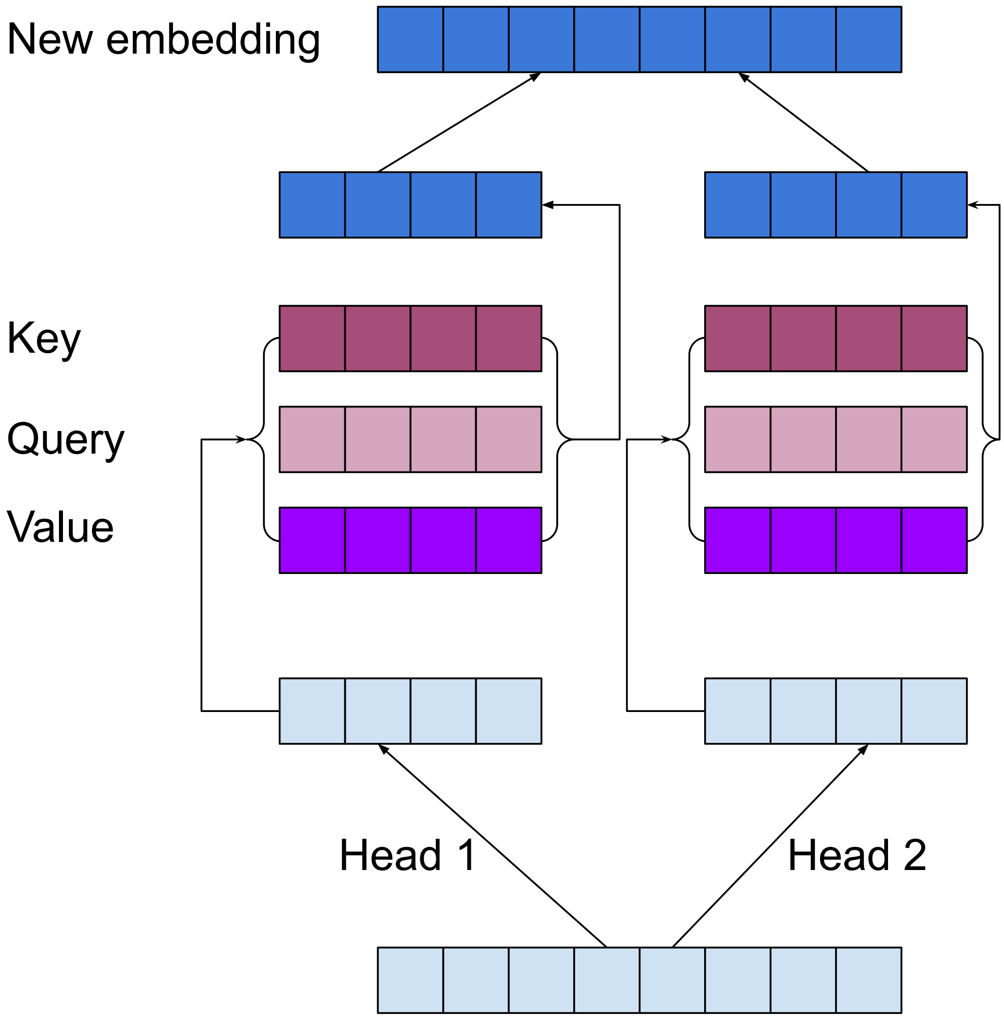

FIGURE IB2: DIAGRAM OF A MULTI-HEAD ATTENTION LAYER

a set of naive embeddings

$$\mathrm{EMB}_0 = [[\mathrm{CLS}], t_1, ..., t_N, [\mathrm{SEP}]]$$

where t_is are the embeddings for all of the tokens, "[CLS]" is a special token that is often treated as the embedding for the text, and "[SEP]" is a special token denoting the end of a piece of text.

As shown in Figure IB2, a multi-head self-attention layer takes in an embedding and outputs another embedding. The input embedding is passed through three linear mappings in parallel to form three matrices: the key matrix, the query matrix, and the value matrix.

$$Q = \mathrm{EMB}_0 W^Q$$
$$K = \mathrm{EMB}_0 W^K$$
$$V = \mathrm{EMB}_0 W^V$$

where Q, K, and V are trainable parameter matrices. Then, for each query, a cosine similarity score is computed between this query and all of the keys, including itself. Then, the value of the token is represented as a linear combination of all of the values of tokens in this piece of text. The weights in the linear combination are the cosine similarities between queries and keys. Mathematically, we have

$$\mathrm{Attention}(\mathrm{EMB}_0) = \mathrm{softmax}(QK^T)V$$

For BERT models specifically, the attention is often calculated as

$$\mathrm{Attention}(\mathrm{EMB}_0) = \mathrm{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where d_k is a scaling factor equal to the number of columns of K. To improve representation capacity, the input embeddings EMB_0 are often broken into multiple equal-sized sub-vectors. The attention is computed for each sub-vector independently and concatenated to output a multi-head attention of the input embedding. In addition, this attention procedure is often repeated many times where the output of the $(i-1)$th attention is normalized and combined with the input of the $(i-1)$th attention to act as the input to the ith attention.

B.3 Bidirectional Encoder Representations from Transformers (BERT)

Similar to the single-layer neural network discussed at the beginning of this section, BERT has an input and an output layer. The input layer takes in a matrix of initial token embeddings, and the output is a set of context-aware embeddings of the same set of tokens. The BERT architecture is an extension of the single-layer network in two ways. First, each building block layer is an attention layer instead of a linear layer. Secondly, BERT has multiple attention layers. As shown in Figure IB3, BERT is an architecture built by stacking multiple encoder blocks. Each encoder block is a sequence of attention layers.

We elaborate on the training objective of BERT, which consists of two parts: Masked Language Modeling (MLM) and Next Sentence Prediction (NSP). During training, a pair of sentences A and B is passed into BERT at a time. A subset of tokens in each piece of text is masked: their token IDs are replaced with the ID of a special "[MASK]" token. After passing through BERT, the embedding for each "[MASK]" token is used to predict the identity of the true token using cross-entropy loss. In addition, the output embedding for the "[CLS]" token is used to predict whether for two sentences A and B, the sentence B is after A. More specifically, this prediction task is created such that there is a 50% chance that sentence B is a random sentence from the training text corpus, and for the other 50% chance, B is actually the sentence following A. The two tasks are trained simultaneously. More specifically, the loss function is a weighted combination of the two losses.

Through this training process, the BERT model is able to produce token and document-level numerical embeddings that represent the context at a token and sentence level. At the end, we can extract the embedding for each token from the last attention layer; in addition, we can extract the embedding of the "[CLS]" token and treat it as the document-level embedding.

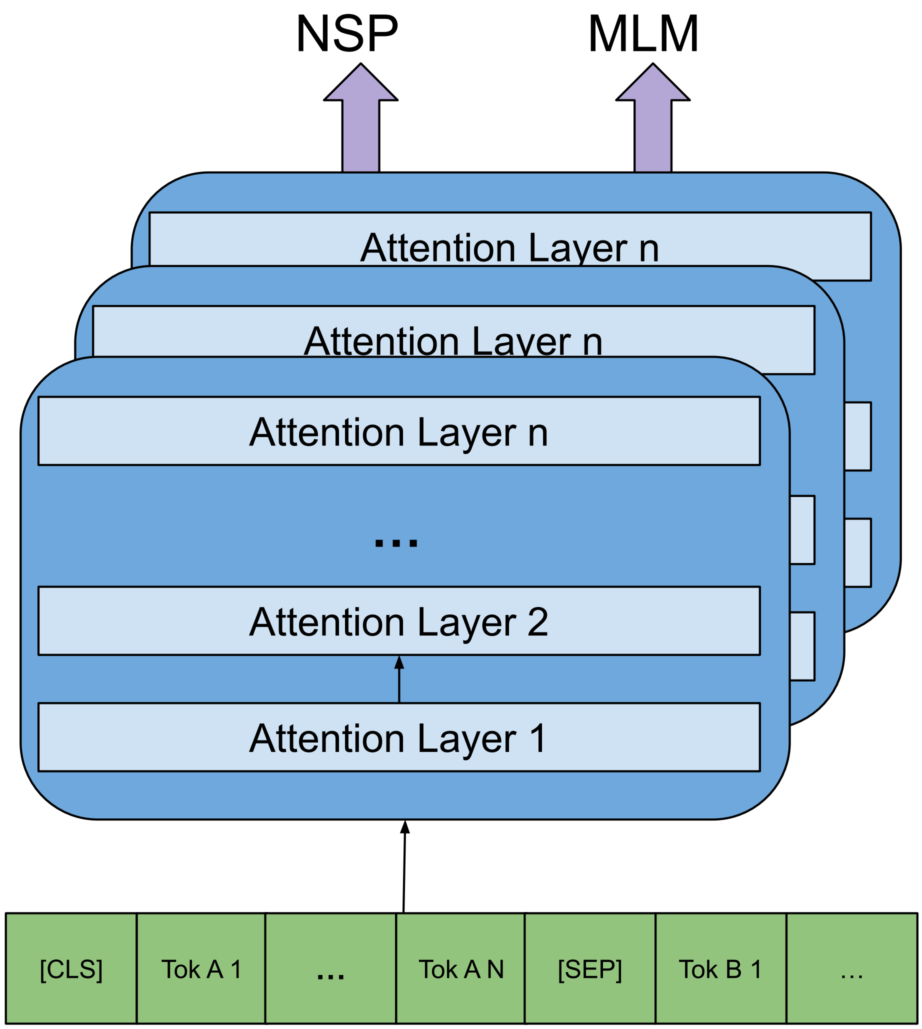

FIGURE IB3: DIAGRAM OF BERT

Appendix C Encoder and I-TEXT

In this section, we describe the encoder architecture as a tool to generate numerical representations of texts. We also describe our I-TEXT method, which utilizes encoders.

C.1 Encoder Architecture

In this subsection, we give an overview of encoder models, which includes the BERT model described in detail in the previous section, and then we describe the specifics of the model we use in our main analysis. The encoder architecture works as follows. Let W denote the original input sentence in words. As shown in Figure IC1, before entering the encoder, W is broken down into three parts: a token embedding E_W^T, which represents the content of the sentence; a segmentation embedding E_W^S, which labels tokens with the sentence they belong to; and a positional embedding E_W^P, which represents the relative distances between each pair of tokens (a "token" is a word or a part of a word if the word is long). A linear combination of these three embeddings then goes into the encoder.

The first step of the encoder is a multi-headed attention layer. Its mechanism can be described as follows. Let E^W denote the input embedding of the encoder. For a given token W_i in sentence W, the embedding is denoted E_i^W. The attention layer calculates the projection of E_i^W onto all token embeddings, including itself, using a dot product. The final output of the single-headed attention layer for each token embedding is a weighted average of all token embeddings, where the weights are the cosine projection coefficient of the current token embedding onto each token embedding. A multi-headed attention layer is analogous to a forest of single-headed attention layers. To construct a k-headed attention layer using a pk-dimensional token embedding, we randomly split the pk-dimensional embedding of each token into k groups of p-dimensional embeddings. We then build a single-headed attention layer with one subset of the token embeddings. Finally, we take a weighted average of all of the outputs of the k heads.

The output of this multi-headed attention layer is then passed through a normalization layer with a residual connection. Residual connection is achieved by passing the input of the multi-headed attention layer directly to the normalization layer, along with the output of the multi-headed attention layer. This residual connection allows gradients to directly flow from the input of the multi-headed attention layer to the next layer while not going through the multi-headed attention layer. After the normalization layer, the output is passed through a feed-forward layer, which converts the output of the normalization layer to the same format as the input of the encoder module. This allows us to stack multiple encoder modules together, where the previous encoder's output can be used as the input for the next encoder. The reason we stack encoders is that the first encoder learns the contextual relationship between pairs of

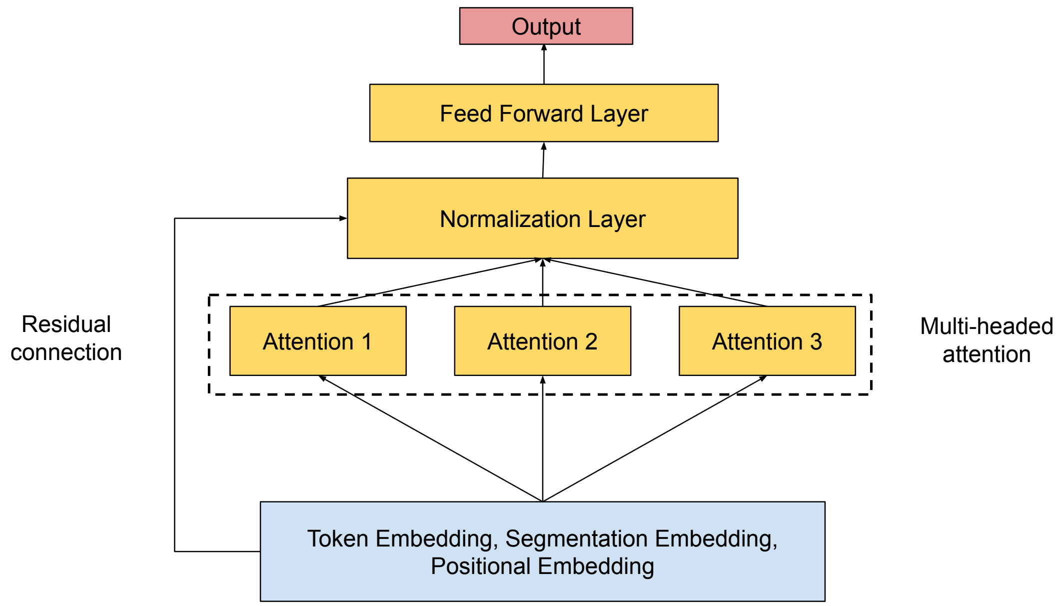

FIGURE IC1: ENCODER MODULE

This figure illustrates the structure of the encoder module. The light blue block at the bottom describes the input. The yellow blocks are the layers within the encoder, and the red block is the output.

tokens, the second encoder learns the relationship between pairs of tokens, and so forth. For the following discussions in this paper, we use the word "embedding" to mean the output embedding of the encoder at the text level.

The pre-trained Longformer model is based on RoBERTa, which uses the encoder architecture to train for masked language modeling (MLM). To train the MLM task, a random subset of tokens in the input sentence is masked with a trivial embedding vector. Then, after this sentence goes through the encoder, the output goes through a fully connected linear layer and a softmax layer to predict what the masked tokens in the original sentence are. This loss is computed using cross-entropy. The final trained RoBERTa model can output embeddings of sentences or texts that represent not only the meaning of the tokens but also the contextual relationship between tokens and sentences.

C.2 Explanation of Inference from Text Analysis

I-TEXT is a neural network-based architecture designed to estimate outcomes associated with a binary variable, where all covariates relevant for analysis are contained within a given text. To use I-TEXT for analyzing the relationship between gender and the recognition of patents, the model must first be fitted. As shown in Figure IC2, the Xs contains four types of information: the texts of patents, the novelty and impact of a patent, the gender indicators of the inventor(s), and the observed number of citations on the patents. There are four neural networks that need to be trained: an encoder model for generating text embeddings, a logit-linear model that maps embeddings along with novelty and impact measures to propensities, and two 2-layer perceptrons that map from embeddings and novelty and impact measures to the estimated number of citations for male and female inventors, respectively. The final loss function is a weighted average of the losses of these four neural networks. After the model is fitted, we can use it to estimate the number of citations of male-written patents if they were written by females and vice versa. As shown in Figure 1, to estimate these outcomes, we run the fitted I-TEXT model where the input data contains the texts of the patents and gender indicators of the inventor(s). The texts are first passed through the encoder model to generate a vector embedding for each patent. The embedding is concatenated with novelty and impact. Then each embedding-gender pair is passed through a classification step: if the inventor(s) are male, the embedding is passed to one network, and if it is written by female(s), the embedding is passed to another network. The resulting number of citations is then estimated by these two networks. In parallel, regardless of the gender indicators, each embedding is passed through the propensity network to estimate the likelihood associated with this patent. Finally, the output of the model is a set of citation-propensity pairs that each correspond to one patent.

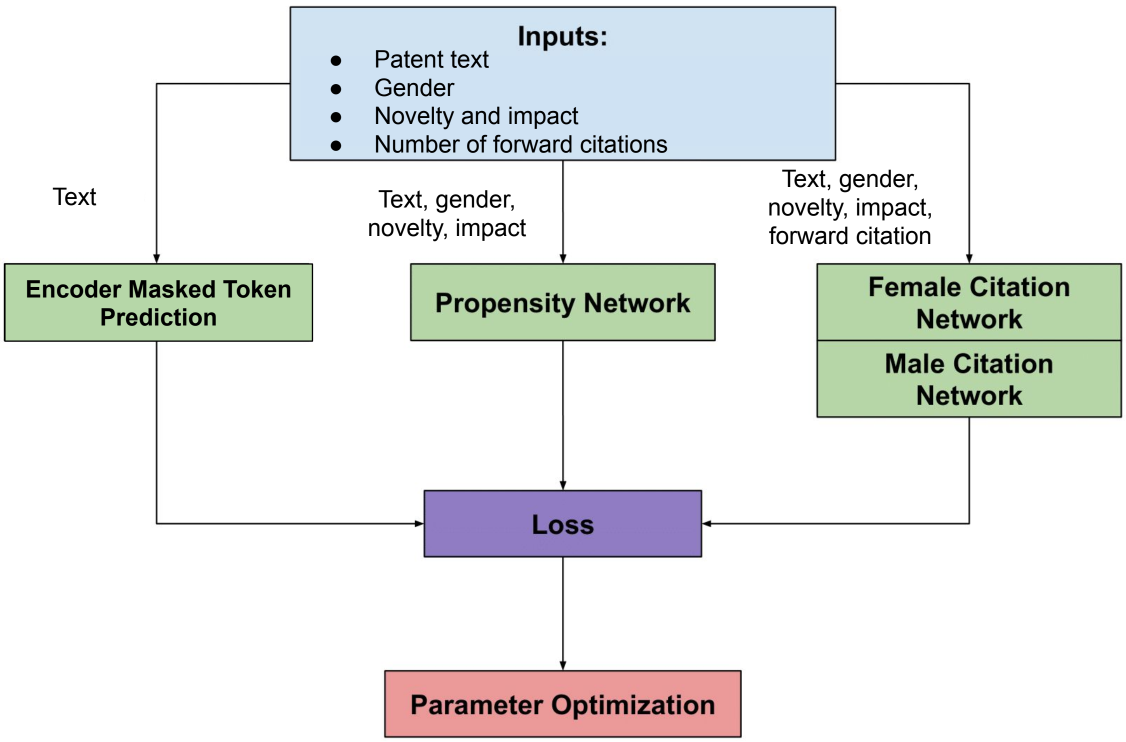

FIGURE IC2: I-TEXT TRAINING PROCEDURE

The figure illustrates the training procedure of I-TEXT once the neural networks are trained. The light blue block at the top describes the input used for estimation. The green blocks are the four neural networks that are trained using the patent data. The purple block denotes the loss function of the model which is a weighted average of the loss of all four networks. Finally, the red block denotes the optimization algorithm that allows the model to get a step toward fitting the training data.

Appendix D Model Fit

An important consideration is the model's fit. First, we discuss how the optimization objective is constructed. We then discuss how well the model performs out-of-sample.

We use a combined loss function for our optimization objective, including the estimation errors of the male and female patent citation neural networks, the propensity network, and the masked language model loss when fine-tuning the text encoder model (Longformer for our main analysis). We estimate our neural networks using 20 epochs. While increasing the number of epochs helps to improve the fitting performance of the neural networks, doing too many runs the risk of overfitting our model to the data. Such overfitting would result in a relatively poor and ungeneralizable representation of patents' text quality. We address this concern by studying the loss function, as presented in Figure IA1, to ensure a reasonable number of training iterations without overfitting the model. Plotting the mean square error (MSE) per batch against the number of passes of the training dataset, we find diminishing improvements to the error rate as we approach 20 epochs, the number of passes utilized in our estimations. This suggests that at 20 epochs, the model has likely at least found a local minimum.

Next, we assess our model's out-of-sample performance by training it for 20 epochs on 80% of the data and then evaluating its accuracy on the remaining 20%. When comparing the estimated log text quality with actual log citations, we obtain a correlation of 0.43, which increases to 0.45 once both measures are exponentiated.

To provide a more intuitive picture of the model's fit, we look at the share of patents where the gap between estimated and actual citations falls below certain thresholds. We find that 68% of patents are off by fewer than 5 citations, 77% are off by fewer than 10, and 90% are within 27. Given that the standard deviation of the test sample is 53, these results point to a solid out-of-sample performance in forecasting patent citations.

Importantly, we note that forward citations in patent data exhibit a heavily right-skewed distribution: while the median patent receives relatively few citations, a small subset of "blockbuster" patents attracts extremely large numbers. These outliers inflate the standard deviation and can create the appearance of low accuracy. Indeed, even standard regression analyses that employ typical controls often leave substantial unexplained variation owing to a handful of exceptionally high-citation patents. An R^2 in the range of 10–15% is common in estimation tasks involving patent citations, reflecting the noisy, skewed, and discrete nature of the data. Our out-of-sample tests confirm that the model captures an economically significant portion of the variation, with 90% of estimated citation counts falling within one standard deviation of actual values.

We stress that our primary objective is not to predict exact citation totals but rather to measure whether a given patent would have received more or fewer citations if its lead inventor were male or female. Even with modest accuracy, the model permits us to detect systematic over- or under-citation relative to the patent's text and observed significance. While a more complex model might replicate underlying biases (given the strong male majority among inventors), our balanced approach preserves interpretability and mitigates overfitting. Consequently, we accept a baseline level of error in exchange for a cleaner estimation of gender differentials. While the large standard deviation reflects the influence of highly cited "superstar" patents, the core inference about undercitation remains robust despite this skewness.

Appendix E Ruling out first mover channel

The timing of patent publication is an important factor contributing to the amount of citations it will receive. In our main analysis, we control for this timing factor using both an importance measure and a time fixed effect. Here we elaborate on this analysis by directly examining if there are any differences between male and female publications in terms of the timing of each publication. More specifically, we analyze whether given two similar patents A and B, a male is proportionally more likely to be the first inventor of the patent published earlier. If males hold the "first-mover advantages" proportionally more often than females, we could expect male-written patents to receive more citations on average. If the reverse is true, we do not expect the timing of publication to positively contribute to the discrepancy between male and female written patents.

We expect the "first-mover advantage" to have the biggest impact on high-impact patents. Therefore, we focus our study on patents in the top 20 percentile in terms of both citations and significance. Similar to our main analysis, we first take a balanced set of patents written by males and females. There are 41730 high-impact patents. We then extract the Longformer embedding for each patent's text and match each patent with the patent that has the most similar embedding vector based on cosine similarity. Within the set of high-impact patents, 17549 have lead inventors whose gender is different from its most similar patent. 45.3% of these patents have female lead inventors. However, out of these 17549 patents, 48.3% of patents were issued before its most similar patent. As shown in Figure IE1, conditional on the lead inventor of a patent being a male, the probability that this patent is published earlier than its most similar high-impact patent (not restricted to the subset of 17549 patents) is 38.5%. Conditional on the lead inventor being a female, the probability is at a higher 43.3%. Therefore, conditional on the proportion of these patents written by females, female lead inventors are more likely to gain a first-mover advantage. This observation should make our estimates in our main analysis conservative estimates.

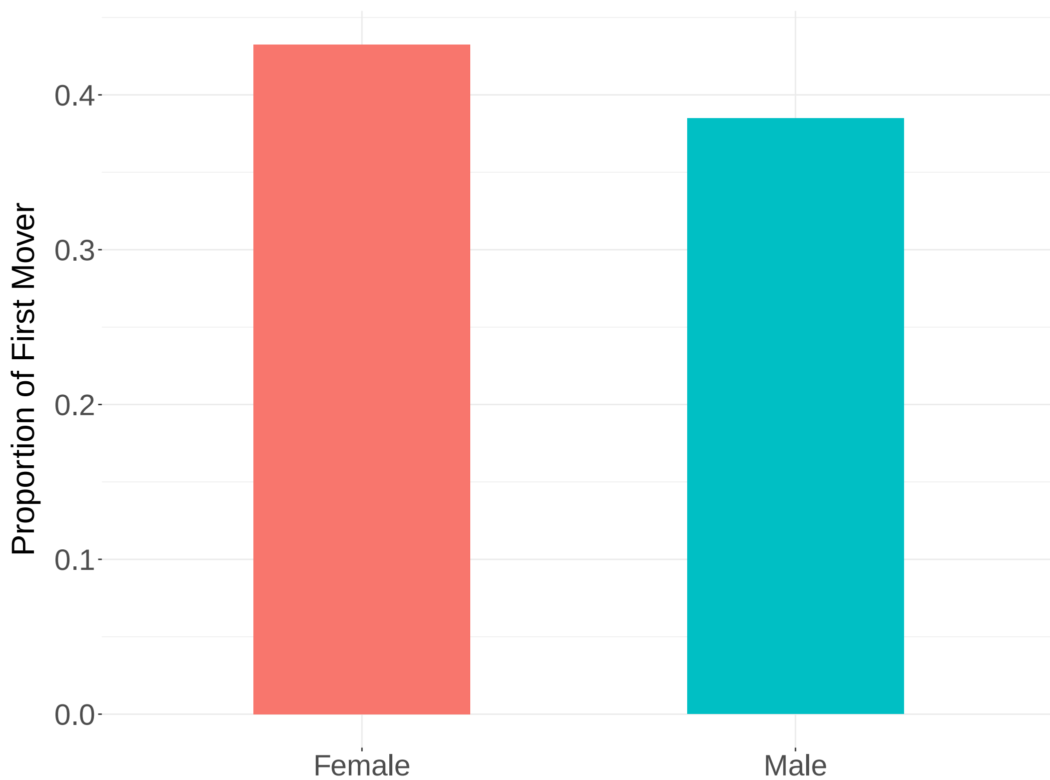

FIGURE IE1: CONDITIONAL PROBABILITY OF BEING THE FIRST-MOVER BY GENDER

This figure plots the conditional probability of the lead inventor's gender to be the inventor of a highly significant patent. The vertical axis measures probabilities. The red left bar measures females, while the green right bar measures males.

Appendix F Word clouds

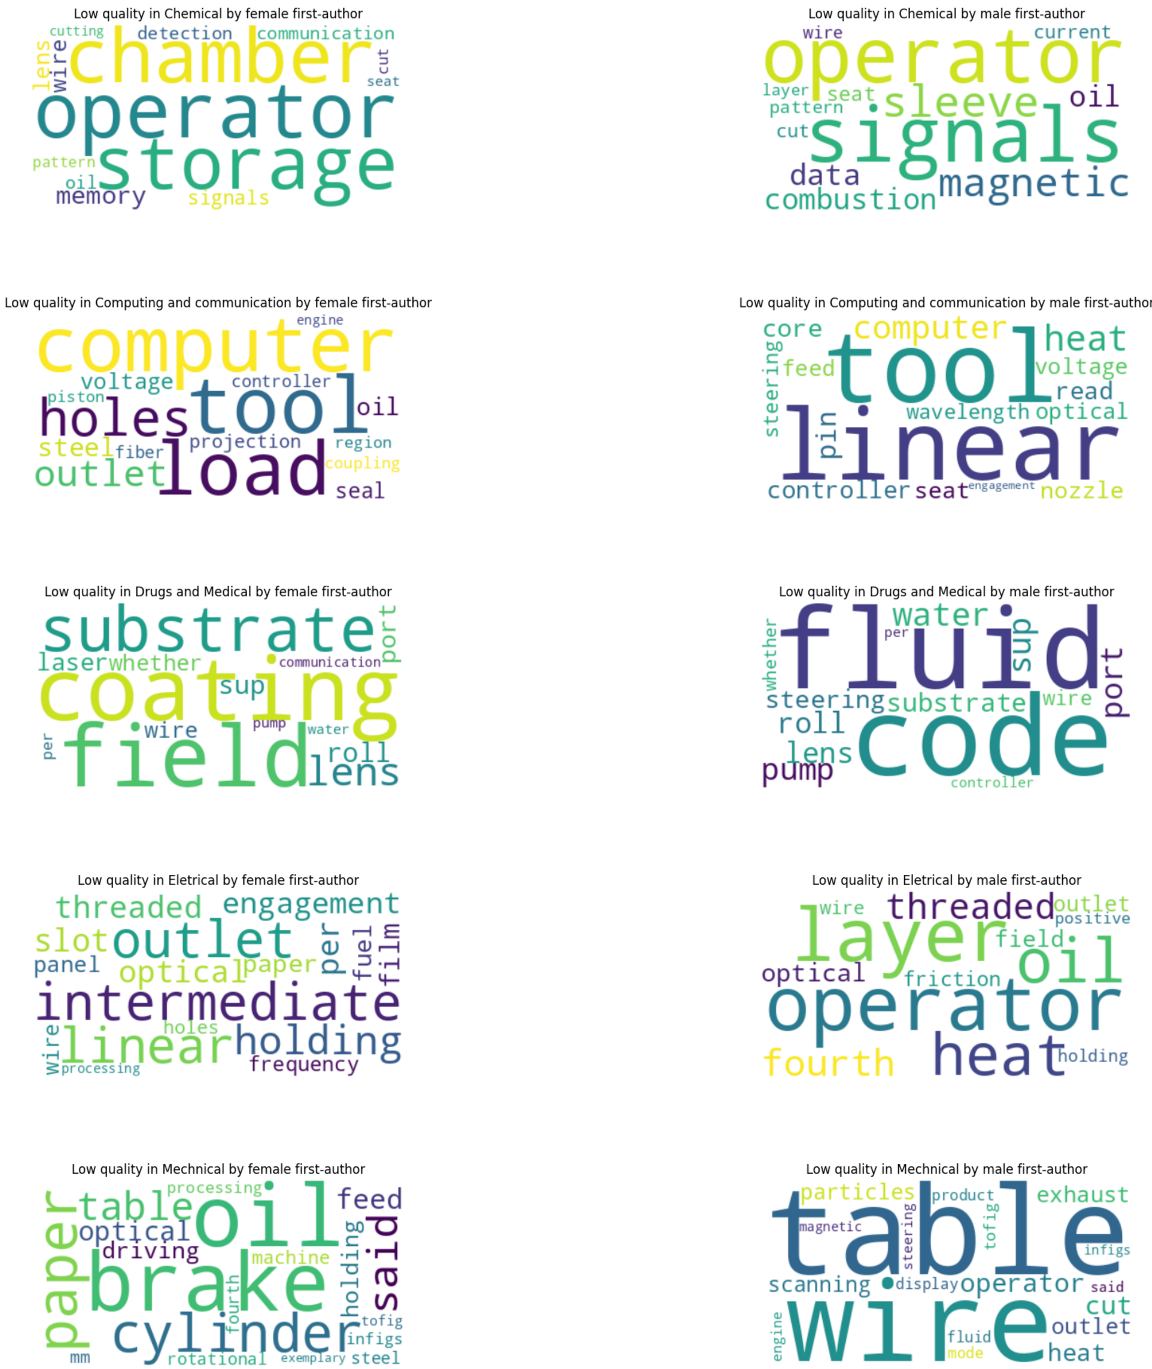

FIGURE IF1: WORDS THAT CORRELATE WITH LOW MODEL ESTIMATED CITATION COUNT (QUALITY)

This figure shows the words that are positively correlated with a low-quality patent in terms of the model estimated citation count in different NBER categories.

FIGURE IF2: WORDS THAT CORRELATE WITH HIGH MODEL ESTIMATED CITATION COUNT (QUALITY)

This figure shows the words that are positively correlated with a high-quality patent in terms of the model estimated citation count in different NBER categories.

Appendix G Additional econometric discussion

Our analysis focuses on extracting the gender-induced bias in the number of citations received by a patent. In particular, we study the bias that comes from the text of the patent and its novelty and impact relative to other patents. In other words, given the same text of a patent and its novelty and impact, how many more or fewer citations would the patent receive if it had a female lead inventor as compared to a male lead inventor? We use the following econometric specification:

$$Y_i(j) = f(Z_i, j) + \beta X_i + \epsilon_{i,j}, \tag{14}$$

where $Y_i(j)$ is the number of citations received or would have received by each patent given j the gender of the lead inventor, $f(Z_i, j)$ is the text-based quality of the ith patent given, Z_i, the concatenated vector that represents the text and its novelty and impact and, j, the lead inventor's gender (0 is male and 1 if female), X_i is a vector of controls like the examiner ID and the art group of patent i, and $\epsilon_{i,j}$ is the i.i.d error term.

In the following list, we describe the steps that we take to estimate $Delta$:

1. **Constructing $\mathbf{Z_i}$**: To produce the vector Z_i, the text W_i is first converted to a numerical embedding through a text encoder and then concatenated with its corresponding novelty and impact to form Z_i.

2. **Estimating text-based quality**: Because Z_i is 770-dimensional (768+2), classic linear models can not fit the data (Ys) well. Therefore, we use two hidden-layer neural networks to model $\hat{f}$ as a piecewise function

$$\hat{f}(Z_i, j) = \begin{cases} \hat{f}_0(Z_i), & \text{if } j = \text{ male} \\ \hat{f}_1(Z_i), & \text{if } j = \text{ female.} \end{cases}$$

   More specifically, we have

$$\begin{aligned} f_0(Z_i) =& \hat{f}_0(Z_i) + \epsilon_i^{f_0} \\ f_1(Z_i) =& \hat{f}_1(Z_i) + \epsilon_i^{f_1}, \end{aligned}$$

   where $\hat{f}_0(Z_i)$ and $\hat{f}_0(Z_i)$ are neural networks estimated using the I-TEXT architecture, and $\epsilon_i^{f_0}$ and $\epsilon_i^{f_1}$ are errors independent of X_i.

3. **Rewriting the citation model ($\mathbf{Y_i}$)**: We combine the error terms of the f estimation and

the error terms of the Y estimation. For brevity, we can write the equation 14 as

$$\begin{aligned} Y_i(j) =& \hat{f}(Z_i, j) + \beta X_i + \left(\epsilon_{i,j} + \epsilon_i^{f_j}\right) \\ =& \hat{f}(Z_i, j) + \beta X_i + \hat{\epsilon}_{i,j}. \end{aligned}$$

4. **Satisfying the overlapping assumption**: We use the other neural network estimated in I-TEXT, the propensity network $g(Z_i) = \text{Prob}(i\text{'s lead inventor is female})$, to remove patents that do not satisfy the overlapping assumption – patents whose text, novelty, and impact are highly suggestive of the gender of its lead inventor.

5. **Computing Delta**: Our *Delta* measure can be written as

$$\begin{aligned} & Y_i(j) - \hat{f}(Z_i, 0) \\ & = \hat{f}(Z_i, j) + \beta X_i + \hat{\epsilon}_{i,j} - \hat{f}(Z_i, 0) \\ & = \mathbb{1}(j = 1)\left(\hat{f}(Z_i, 1) - \hat{f}(Z_i, 0)\right) + \beta X_i + \hat{\epsilon}_{i,j}. \end{aligned}$$

Finally, by running a linear regression with controls X_i, we can recover the coefficient on the gender indicator which is the gender-induced bias from texts.

Appendix H Sample Selection

A key detail to your paper is how our subsamples are constructed. We tailor our analysis by constructing specific subsamples from the data. The subsamples are constructed using the following workflow. First, we merge the entire set of patent texts with abstracts, the gender of the first author, and the number of forward citations. Then, it is merged with a larger data set that contains the patents' texts. The initial merge is different for different subsamples. For example, for the male inventors study, only patents with at least one citation coming from a male inventor are included in the merge.[22] After the initial merge, a balanced subsample of male-led patents and female-led patents is taken from the overall sample. This is because part of our model fits a neural network that estimates the propensity that a given patent has a female first author. Without the subsampling, over 90% of the patents (depending on the specific study) are written by male first authors. Therefore, the network will likely have a bias towards labeling everything as male-authored. Taking a balanced subsample alleviates this concern.

The next part of our analysis removes the patents that have extreme propensity estimates. As discussed in section 2, in order to satisfy the overlapping assumption of inference, we drop the patents that are either estimated to be highly likely to be male-written ($g(z) < 0.03$) or female-written $g(z) > 0.97$. The remaining dataset is used in the regression analysis.

The list below shows all of the tests where the neural network is refitted and the corresponding number of observations used in the regression to compute the *Delta* measure after creating the balanced subset, dropping missing values in independent variables, and dropping data points with extreme propensity scores, below 3% or above 97%.:

All samples are balanced to have the sample number of patents by lead female and male authors, unless otherwise noted. All models are trained with the *Significance* measure from Kelly, Papanikolaou, Seru, and Taddy (2021) in the neural network unless otherwise stated.

Appendix I Additional Training Details

We begin by mapping the embeddings generated by the language model and additional controls, such as novelty and impact, to two separate 200-dimensional fully connected layers, each corresponding to the male and female citation estimation components of the model. Both layers employ the ReLU activation function. Simultaneously, the same embeddings are passed through a sigmoid activation to produce a propensity score, representing the estimated probability that a given patent is authored by a female lead.

[22] The significance measure from Kelly, Papanikolaou, Seru, and Taddy (2021) is also merged onto the dataset for our main baseline analysis and examiner-inventor analysis in the main body of the paper. This merge remove any patent post 2015 because the latest patent with a non-empty 5-year significance measure is in 2015.

The loss function comprises three components: (i) the sum of the mean squared error (MSE) losses from the male and female citation estimators (estimated log citations versus actual log citations); (ii) the binary cross-entropy loss associated with the gender propensity estimation; and (iii) the masked language modeling (MLM) loss for predicting the identity of a single masked token in each input sequence of length 512. The same loss function setup is used for all different models and architectures in this paper. The overall loss is computed as the sum of these three components, each assigned an equal weight, $(1, 1, 1)$.

Model training is performed for 20 epochs with a 10% warmup period. We employ the AdamW optimizer, using a learning rate of 2×10^{-5} and otherwise default parameter settings, following standard practice in the fine-tuning of large language models. Training proceeds in batches of 16 observations.[23]

[23] We are unable to fit the model on a V100 GPU with a maximum sequence length of 1024 even with a small batch size like 16. With our current hyperparameters and balanced subsample, training requires more than *three weeks* to complete.

TABLE IH1: SAMPLE CONSTRUCTION DETAILS

Description	Corresponding table numbers	# Observations	Significance in neural network?
Positive-citation patents	Table 1; Table 2; Table IA1	521,967	In NN
Positive-citation patents, emerging fields	Table 3	494,250	In NN
Positive-citation patents, male examiners	Table 4	72,505	In NN
Positive-citation patents, female examiners	Table 4	17,603	In NN
Positive-citation patents, male inventors	Table 5	266,015	In NN
Positive-citation patents, female inventors	Table 5	4,727	In NN
Positive-citation patents with value data	Table 6	191,075	In NN
All patents (full sample)	Table IA1	3,889,331	In NN
Positive-citation patents	Table IA2	554,486	Not in NN
All patents	Table IA2	738,386	Not in NN
Positive-citation patents; female-authored = solo-female	Table IA3	193,725	Not in NN
Positive-citation patents; SciBERT props/expectations	Table IA4	654,552	Not in NN
Placebo: random gender (50/50) on positive-citation sample (5% subsample)	Table IA5	204,695	In NN
Citations window: 0–1y	Table IA6	11,637	Not in NN
Citations window: 2–5y	Table IA6	244,104	Not in NN
Citations window: 6–10y	Table IA6	259,436	Not in NN
Citations window: 11–20y	Table IA6	308,207	Not in NN
Propensity cutoffs 5/95 (vs 3/97)	Table IA7	472,332	In NN
Positive-citation patents with writing-quality metrics	Table IA8	553,573	In NN
Writing-quality + controls	Table IA9	553,569	Not in NN
Define Δ using female model	Table IA10	521,967	In NN
Define Δ using female model	Table IA10	554,486	Not in NN
Include authors' prior counts in NN	Table IA11	550,198	In NN
Exclude self-citations	Table IA12	462,717	In NN
DV: asinh(Δ)	Table IA13	554,486	Not in NN
No propensity trimming	Table IA14	687,830	In NN
Public-firm positive-citation patents	Table IA15	191,086	In NN
Expected citations as regressor (Iaria et al., 2024)	Table IA16	521,967	In NN