Inventor Gender and Patent Undercitation: Evidence from Causal Text Estimation[†]

Yael V. Hochberg Rice University & NBER

Ali Kakhbod UC Berkeley

Peiyao Li UC Berkeley

Kunal Sachdeva Rice University

August 2023

Abstract

Implementing a state-of-the-art machine learning technique for causal identification from text data (C-TEXT), we document that patents authored by female inventors are under-cited relative to those authored by males. Relative to what the same patent would be predicted to receive had the lead inventor instead been male, patents with a female lead inventor receive 10% fewer citations. Patents with male lead inventors tend to undercite past patents with female lead inventors, while patent examiners of both genders appear to be more even-handed in the citations they add to patent applications. For female inventors, market-based measures of patent value load significantly on the citation counts that would be predicted by C-TEXT, but do not load significantly on actual forward citations. The under-recognition of female-authored patents likely has implications for the allocation of talent in the economy.

JEL classification: J16, J24, J71, O30, C13

Keywords: Innovation, Gender, Patent, Machine Learning, Big Data, Causal Inference

[†]We thank Tania Babina, Daniel Bergstresser, Po-Hsuan Hsu, Amir Kermani, David Levine, Zack Liu, Dmitry Livdan, Ulrike Malmendier, Asaf Manela, Danqing Mei, S. Lakshmi Naaraayanan, 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), and Boston University (TPRI IP Day) for their helpful comments and suggestions. E-mail: Hochberg (hochberg@rice.edu), Kakhbod (akakhbod@berkeley.edu), Li (ojhfklsjhl@berkeley.edu), Sachdeva (kunal.sachdeva@rice.edu).

Female inventors appear to face significant obstacles when seeking patents. Women face significant 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 more generally (Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019); Reshef, Aneja, and Subramani (2021)). 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 female-authored patents in fact receive fewer forward citations relative to male-authored patents (Jensen, Kovács, and Sorenson (2018)), suggesting either that patents granted to female inventors are of lower quality, or, more concernedly, that the quality of their patents are not fully recognized in the form of forward citations, a commonly used measure.

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 causal contribution of gender to the citation of patents of similar quality. Our methodology builds on a burgeoning set of research in the computer science literature that studies causal identification using textual data (see e.g. Khetan et al. (2022); Shao et al. (2021)). The intuition behind these models is straightforward. Our goal is to identify the expected change in outcome if we apply treatment while holding fixed any mediating variables affected by the treatment that also might affect the outcome.

Our approach, which we label as Causal Text Analysis, or C-TEXT, estimates causal effects from observational text data, adjusting for confounding features of the text, such as the subject or writing quality. It assumes that the text content suffices for causal identification but is prohibitively complex for standard analysis. Our C-TEXT approach utilizes causally sufficient embeddings, relatively low-dimensional document representations that preserve sufficient information for causal identification, thus enabling efficient estimation of causal effects. The causal sufficiency reduces dimensionality yet preserves aspects of the text that predict both the treatment and the outcome while disposing of linguistically irrelevant information–which is also causally irrelevant. The identification assumption is that the text contains all unobserved information necessary to measure the desired effects (quality of the patent and forward citations, conditional on gender). Our C-TEXT model generalizes embedding-specific approaches such as Veitch et al. (2020) and allows for various encoder architectures (here, Longformer and SciBERT) to causally identify treatment effects.

We then use the resulting embeddings to train two neural networks—one per gender of the lead patent author—using the embedding's numerical representations of the patents as inputs, and forward citations as outputs. Each neural network represents a mapping from embedding vectors to citation counts. The first mapping is trained using the subset of data where the patent is female-authored, while the second mapping is trained using the data where the patent is male-authored. Unlike the standard OLS approach, the neural network approach captures complex and often nonlinear relationships between inputs and outputs, particularly when dealing with high-dimensional inputs.

Having obtained parameters for each gender's citation-prediction model, we then take the sample of patent data and run it through the citation model trained on its own gender's inputs and outputs and through the citation model trained on the *opposite* gender's inputs and outputs. This produces a set of counterfactual citation counts for each patent, holding all else equal and changing only the gender of the authors.[1]

Our main sample covers all utility patents granted by the U.S. Patent Office (USPTO) from 1976 through 2021. For our main analyses, we focus on the first inventor's name

[1]The methodology also incorporates a gender propensity model to ensure a patent's text is not identified as male- or female-authored, to ensure a quality counterfactual can be computed. Dropping this restriction only strengthens our results.

on the patent (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-authored patents, including restricting to single-authored patents or majority-gender teams.

We begin by documenting that even with no adjustments for patent quality or characteristics, there is a statistically significant difference in the number of forward citations for patents with a female first author versus those with a male first author in our matched data, consistent with Jensen, Kovács, and Sorenson (2018). Patents authored by women appear to be less likely to receive any citations than male-authored patents. This pattern persists when we control for factors such as the identity of the patent examiner, the art unit of the patent, the identity of the attorney who assisted with the patent preparation and submission, the assignee, and the patent issue year.

Next, we use the C-TEXT methodology to mediate the differences in the quality of the patents in order to identify the causal effects of gender on forward citations. Our analysis then proceeds as follows. First, we define two primary datasets upon which we run our tests. In the extensive margin sample, we define the analysis sample as all patents, including those that receive zero citations (the modal patent). In the intensive margin sample, we restrict to those patents that receive at least one forward citation. Next, we apply the C-TEXT methodology to obtain counterfactuals for the opposite gender. Using these counterfactuals, we calculate the average treatment effect (ATE), the difference between the predicted forward citations for the patent out of each gender's neural network, averaged over the treated sample. In addition, for each patent, we then calculate the *Delta* for each patent, which we define as the difference between the *actual* number of forward citations received and the predicted forward citations had the author(s) been male. Finally, we use *Delta* as the dependent variable in simple OLS regressions that allow us to include a variety of controls and fixed effects for added robustness.

Our baseline estimates suggest that patents with a female first author would have received more citations if their first author had been male. At the extensive margin, we find that the average treatment effect for the sample is -1.95. In other words, for any given patent, our trained models would estimate nearly 2 fewer forward citations if the first inventor listed was female than if the first inventor listed was male, holding the patent content and writing equally. Applying the C-TEXT methodology and comparing actual citations to those that would be predicted if the patent was authored by a male lead inventor, we find that patents with a female lead inventor received approximately 13.5% fewer citations than an equivalent quality patent in the same art unit, evaluated by the same examiner, would receive had the lead inventor been male. This difference equates to approximately 1.9 fewer citations per patent. The impact of this undercitation is most pronounced for the most impactful patents, with female lead-authored patents being less likely to reach the top decile of citations. At the intensive margin, we find similar effects, as the ATE is -2.66. In OLS regressions, patents with a female lead inventor received approximately 8.7% fewer citations than an equivalent patent would receive had the lead inventor been male, a difference again of approximately 1.86 fewer citations per patent.

The results are robust to various alternative specifications and are not attributable to sample selection or model overfitting. The results are also robust to various approaches to defining a "female-authored" patent. For example, we obtain similar results when comparing patents with a single female author to those with a single male author or patents with author teams composed of a majority of female authors versus patents authored by a majority of male authors.

Our results hold across Cooperative Patent Classification (CPC) major categories and subcategories of patent technology, with heterogeneity by subcategories. We observe similar patterns when using the National Bureau of Economics (NBER) classification system. We observe that the undercitation is particularly large in emerging technology fields. This undercitation of female-authored patents has grown over time, becoming pronounced through the year 2000, with it recently shrinking in the 2010s.

Identifying the *source* of this undercitation is clearly of interest. The citations inventors themselves could drive undercitation of female-authored patents include in their patent applications, or by citations that are added by patent examiners, and may depend on the gender of the examiner or inventor adding the citation. Controlling for art unit, issue year, examiner, attorney, and assignee, patents with male first authors significantly undercite patents with a female first author. In contrast, we see little evidence of female first authors, or either male and female examiners underciting female-led patents. Overall, the results suggest that the undercitation of female patents is largely due to patents with male first authors underciting past female-authored patents in their patent applications.

Finally, we explore whether the economic value of patents, as calculated by the market reaction to its issuance, correlates more with actual citations or with the predicted citations for the same patent if male lead authored based on our methodology. We use the economic value of a patent as measured by public markets from Kogan, Papanikolaou, Seru, and Stoffman (2017), which is generally considered to be forward-looking and determined at the time of issuance. For all female-authored patents, we regress these measures of economic value on actual forward citations and the forward citations that would be predicted for the same patent if male-authored. When horse-raced against each other, measures of expected economic value load significantly on the predicted forward citation measure, but do not load significantly on actual forward citations. The results suggest that expected economic value, as measured by market reactions, maybe a less biased proxy for patent quality than standard measures of realized forward citations.

Our results come with several important considerations. First, causal text analysis relies on the assumption that the text analyzed captures all the unobserved factors that should influence the outcome being examined. While it is not possible to test this assumption directly, it is reasonable to assume that the content of the patent is closely related to its quality or importance, the unobserved factor of interest. Second, assessing the goodness-of-fit of any given model in computing counterfactual outcomes is challenging. We utilize a number of approaches to assess model fit. Finally, although our evidence suggests women receive fewer citations for patents of equal quality, we do not argue that this represents discrimination, as we cannot observe the intent of examiners or inventors. Further research will be necessary to establish *why* patents with female lead inventors are undercited.

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).[2] 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), 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 female-authored patents are systematically undercited relative to their actual quality, the use of 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 contribute to the literature on the impediments that women and minorities face in obtaining patents, with emphasis on the unequal application of laws (Cook (2014)), unequal opportunities (Cook (2020); Cook and Kongcharoen (2010)), and discrimination by patent examiners (Desai (2019)). These obstacles result in depressed levels of applications and lower success rates for females in obtaining patents (Jensen, Kovács, and Sorenson (2018)). In contrast to this literature, which focuses on identifying differences in patent applications and approvals, our findings focus on a relatively unexplored question: whether women also face obstacles in *citation* to their patents.

Our findings also contribute to the broad literature studying obstacles that women face in various research fields. Recent work by Sherman and Tookes (2022) documents that women face discrimination in financial economics publishing and job placement. Sarsons, Gërxhani, Reuben, and Schram (2021) and Sarsons (2017) show women receive less credit attribution for co-authored work in economics, while Hengel and Moon (2020) show that, controlling for quality, male economist receive fewer citations for their work in the "top-five" journals. Related, Card, DellaVigna, Funk, and Iriberri (2020) shows

[2]In related research, the marginal investor values patents Aghion et al. (2013); Hall et al. (2005); Hirschey and Richardson (2004); Hirshleifer et al. (2013).

that female-authored papers receive about 25% more citations than observably similar male-authored papers, while Koffi (2021b) finds that undercitation in economics is more likely to be of women-authored papers and that male authors are more likely to cite male-authored papers. Koffi (2021a) find that female-authored economics papers are more likely to be cited outside economics, less likely to be cited by top-tier journals, and less likely to be cited by men. Chawla (2016) and Koffi and Marx (2023) study broader academic fields. Our work suggests parallels in patent citations as well.

In addition, our paper makes an important methodological contribution. In economics, a rapidly growing branch of the big data literature uses natural language processing to quantify text (see e.g. Gentzkow, Kelly, and Taddy (2019a)).[3] Our paper introduces new methods, based on recent advances in computer science, that allow the use of text embedding to mediate and identify *causal* effects to the economics literature. To the best of our knowledge, we are among the first researchers to apply deep learning in economics for causal inference using language.

1 Data

Our main analysis uses data on patent content, citations, and attributes. Our main sample covers all utility patents the U.S. Patent Office (USPTO) granted from 1976 through 2021.

1.1 Patent Content

Our sample of patents comes from the USPTO's Patent Examination Research Database (PatEx) dataset. In our main analyses, We study the quality of the patents through the lens of patent text, as they should provide a clear summary of the core contribution of the patent. Importantly, this is the key text input into the C-TEXT model. In robustness

[3]A partial list of papers in this vein includes the work of Athey and Imbens (2019); Bellstam, Bhagat, and Cookson (2021); Cong, Liang, and Zhang (2019); Erel, Stern, Tan, and Weisbach (2021); Gentzkow, Kelly, and Taddy (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).

tests, however, we also consider SciBERT, a fine-tuned text embedding model specifically designed for scientific writing. Because SciBERT is only computationally feasible for shorter texts, we utilize patent abstracts for this analysis. Our results remain qualitatively similar.

1.2 Gender of Inventors

Our main treatment variable is the gender of the lead inventor (first author).[4] The person who is named first on a patent is usually the primary contributor. Moreover, the first name listed on the patent may be more salient, 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.

One challenge when studying gender and patents is that women are underrepresented as inventors on patents (Hunt et al., 2013). As a result, in order to avoid discrepancies in the predictive power of the male and female-trained neural networks, we must first balance our sample across patents with lead inventors from each gender. To do this, we first use all patents with a female lead inventor and extract a random subsample of patents with male lead inventors of the same size.

In some cases, women and men can differ substantially in writing style. To ensure there is some level of ambiguity as to whether it was authored by a male or female, we follow the approach in Veitch, Sridhar, and Blei (2020), and estimate a propensity model using a one layer logit-linear neural network, where the objective function is the binary-cross-entropy between the predicted treatment indicator and the true treatment indicator. Using the text of the patent, the output of this neural network is the predicted probability that a female lead author writes this patent. We then drop (i) all patents in the male subsample whose estimated propensity of being female-authored based on the text is very low (less than 3%) and (ii) all patents in the female subsample whose estimated propensity of being female-written based on the text is very high (greater than 97%).[5] This step results in the variation in the sample sizes across our subsample

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

[5]All our results remain qualitatively similar in nature and stronger in magnitude if we do not exclude

tests.

1.3 Patent Citations

Patent forward citation counts are obtained through the use of 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 female-authored patents. We therefore distinguish between patent forward citations (the easily observed outcome for a patent) and quality, which is the measure of true interest and can be mediated by using the content of the patent.

Patent forward citations are highly skewed in their distribution, with only a few patents receiving a disproportionately high number of citations. As an alternative to simple counts of forward citations, we also consider whether a patent receives citations in the top decile of all patents.

1.4 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 that are 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 only record citations that examiners and inventors explicitly added.

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

patents whose author gender can be clearly identified from the text content alone.

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 examiner of each patent. We rely on the name of the first lead examiner for patents with multiple examiners due to their prominence. 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). We assign a gender when the percentage of names in the state belonging to that gender is above 70%.[6] 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).

1.5 Other Patent Attributes

When an inventor files a patent application with the USPTO, the application is assigned a USPC class and subclass based on its field of technology. The application is then assigned to an "art unit" comprised of several examiners who specialize 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 898 art units and 11,953 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.[7]

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

Descriptive statistics for our sample are presented in Table 1.

[6]We take a conservative approach and apply a high confidence interval to reduce Type I errors when identifying males and females.

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

2 Empirical Strategy

Our 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 compute the counterfactual of citations based on the gender of the inventor.

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

2.1 High-Dimensional Representation of Patent Text

There are a variety of possible approaches to transform text into numerical form. Many of these neural embedding models stem from the seminal Bidirectional Encoder Representations from Transformers (BERT) model which transform each piece of text into a high-dimensional numerical vector. Developed by Google (Devlin, Chang, Lee, and Toutanova, 2018), BERT has become the leading approach in many commercial applications, including Google's search platform. 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)).[8]

For our main results, we use an encoder, Longformer, which is a special version of a "Robustly Optimized BERT Pretraining Approach" (RoBERTa) adopted to excel at representing longer texts using newly developed mechanisms like global and local attention (Beltagy, Peters, and Cohan, 2020). Its base model, RoBERTa, expands and refines the BERT architecture. In contrast to BERT, RoBERTa is only trained on masked token prediction, is trained using a larger sample that contains longer texts, and is trained over a

[8]See Jha, Liu, and Manela (2022) for an excellent discussion of BERT.

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). In our application, the encoder model produces a high-dimensional representation with 768-dimensional embeddings to represent the text of the patent. We describe the encoder architecture in detail in Appendix B.

2.2 Causal Text Analysis (C-TEXT)

Having created high-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 leading machine learning technique called Causal Text Analysis (*C-TEXT*) that allows us to causally estimate the contribution of language on a binary treatment variable, using the text of the patent as a mediator for quality.[9] C-TEXT comes from recent advances in computer science, including Khetan, Ramnani, Anand, Sengupta, and Fano (2022); Shao, Li, Gu, Qian, and Zhou (2021); Veitch, Sridhar, and Blei (2020). Causal text analysis allows us to use the text of patents as a mediator to causally identify the role of gender on patent citations (Figure 1). To the best of our knowledge, ours is among the first papers in economics to apply deep learning to causal inference with language.

C-TEXT is a neural-network-based architecture that estimates counterfactuals of a binary treatment under the assumption that all of the unobserved information needed for causal identification is contained within a given text. As shown in Figure IB1, the input data for training contains three types of information: the text of the patents, 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: a Longformer model for generating text embeddings, a logit-linear model that maps embeddings to treatment propensities, and 2 two-layer perceptrons that map from embeddings to male and female predicted number of citations, respectively. The final loss function is a weighted average of the losses of these four neural networks.

[9] See Appendix A for a more indepth discussion.

The C-TEXT methodology has two key steps. First, it uses a language model. Given the context of our question, this paper uses Longformer, a transformer-based model to transform the text of each patent into a high-dimensional numerical vector. The embedding vectors are numerical representations of the text that looks to preserve both the meaning of individual words and the underlying context of each word. Second, C-TEXT estimates the number of citations an inventor would have received if that person were assigned the opposite gender. This is accomplished by training two neural networks, where each model represents a mapping from embedding vectors to our outcome variable, forward citations, with the first mapping trained using the subset of patents with a female lead inventor and the second mapping trained using the subset of patents with a male lead inventor. The two estimated mappings, combined with the high predictive performance of neural networks, allow us to approximate the true mappings.

Armed with our two mappings, we can then estimate the counterfactual of gender on citation. That is, 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?

The procedure is depicted in Figure 2. First, we run the trained C-TEXT model where the input data contains the texts of the patents and gender indicators of the author(s). The texts are first passed through the trained Longformer model to generate a vector embedding for each patent. Then each embedding-gender pair is passed through a decision step: if the author(s) are male, the embedding is passed to the female citations network, and, if the author(s) are female, it is passed to the male citation network. The counterfactual number of citations is then computed by these two networks. In parallel, regardless of the gender indicators, each embedding is passed through the propensity network to estimate the treatment propensity of this patent, which is used to identify patents that are clearly predicted (97%+ probability) to have been written by one gender or the other irrespective of quality, which are then dropped (as discussed in section 1). Finally, the output of the model is a set of counterfactual citation-treatment propensity pairs that each correspond to one patent.[10]

[10]The citations estimates using the true gender network is also saved for computing the ATE and ATT.

The framework can be expressed more formally in mathematical terms. We denote the text in the abstract of the ith patent as W_i. We fine-tune the Longformer model f to map W_i to Z_i where Z_i is the embedding of the patent text. Then we use a logit-linear network g to map Z_i to a real number, which represents the treatment propensity of this patent. Here the treatment propensities are the probability that this patent has a female lead inventor.

$$g(Z_i) = P(T_i = 1|Z_i) = (g \circ f)(W_i) \tag{1}$$

In addition, we have two citation networks Q_1 and Q_0. Q_1 maps an embedding vector to the predicted number of citations if the patent has a female lead inventor, and Q_0 maps an embedding vector to the predicted number of citations if the patent has a male lead inventor. Mathematically, we define a piecewise mapping Q that represents the two networks:

$$Q(T_i, Z_i) = \mathbb{E}(Y_i(T_i)|Z_i) = \mathbb{E}(Y_i(T_i)|f(W_i)) \tag{2}$$

where $Y_i(0)$ and $Y_i(1)$ denote the potential outcomes of the ith patent. In our case, these potential outcomes are the number of forward citations. Given these mappings represented by neural networks, we can then estimate the average treatment effect (ATE) and the average treatment effect on the treated (ATT) using the following equations for a set of N patents.

$$\text{ATE} = \frac{1}{N}\sum_{i=1}^{N}\left[\mathbb{E}(Y_i(1)|Z_i) - \mathbb{E}(Y_i(0)|Z_i)\right] = \frac{1}{N}\sum_{i=1}^{N}\left[Q_i(1, Z_i) - Q_i(0, Z_i)\right] \tag{3}$$

$$\text{ATT} = \frac{1}{\sum_{i=1}^{N} T_i}\sum_{i=1}^{N} T_i\left[\mathbb{E}(Y_i(1)|Z_i) - \mathbb{E}(Y_i(0)|Z_i)\right] = \frac{1}{\sum_{i=1}^{N} T_i}\sum_{i=1}^{N} T_i\left[Q_i(1, Z_i) - Q_i(0, Z_i)\right] \tag{4}$$

While the model mediates for the quality of the patent based on text, 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 assignee, which are also not included in the patent text. As a result, to further refine the measured ATT. We next pass the difference between actual forward citations and predicted forward citations in the absence of treatment through OLS regressions with fixed effects for the above. 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, and the actual gender of the leader inventor (treatment 1/0, where 1 is female, and 0 is male), and other patent characteristics.

2.3 Assessing C-TEXT's Identification Assumptions

There are three assumptions for C-TEXT that the econometrician must consider. We discuss each of them and how they are satisfied in our setting.

2.3.1 Text Renders the Effect Identifiable

The first necessary condition is that the text of the documents must render the effect identifiable. Said differently, the effect that the econometrician is measuring must be measurable directly from the text. Similar to an exclusion restriction within other identification strategies, this cannot be formally tested. Instead, this condition must be inspected and potentially falsified by considering other channels.

In the context of this paper, the effect we wish to measure is the relationship between quality and forward citations in the presence or absence of treatment. Our identifying assumption is that the quality of the patent should be measurable by the content (text) of the patent itself. Patent examiners read the patent application's text to evaluate the patents' novelty before granting a patent. 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 the prediction of both treatment, T, and outcome, Y. In our setting, this means that embedding, a lower-dimensional representation of the text, is sufficient to capture the gender and quality of citations.

In order to evaluate the efficacy of our embedding representations, Longformer, we utilize synthetic tests to determine the precision of our model. This process begins with calculating synthetic results for every patent in the comprehensive dataset. To achieve this, we employ a randomly generated linear transformation. This transformation utilizes a uniformly random 768-element vector with values ranging from 0 to 1. Following this, we compute the dot product of this random vector with each patent's 768-dimensional embedding. The calculated values represent the synthetic outcomes for both female. The male synthetic outcomes are created by shifting each female synthetic outcome by a known true treatment effect. As we know the true treatment effect, this model allows us to evaluate the model's performance effectively. When the C-TEXT model is applied, it successfully reveals the known true treatment effect with a high level of accuracy, suggesting that Longformer is proficient in extracting semantically significant information from text.

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 be consistent. That is, the treatment and control groups should have common support.

To address this, as discussed above, we follow the procedure of Veitch, Sridhar, and Blei (2020) and drop the patents with either below 3% treatment propensity or above 97% treatment propensity. In our study, the treatment is the female gender indicator of the lead inventor. Therefore a treatment propensity of at most 3% implies that this patent, as defined by the embedding of the text, almost certainly has a male lead inventor. On the other hand, a treatment propensity of at least 97% implies this patent almost certainly has a female lead inventor. This procedure preserves over 80% of our data after dropping the propensity score outliers. Importantly, our results remain robust, suggesting that the conditional outcome and propensity score models are consistent.

3 Do Citations to Patents Differ By Gender of Inventor?

Do forward citation counts for patents differ across 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 C-TEXT implied ATE and ATT based on patent text alone. We then utilize the C-TEXT output to further refine the estimates of causal differences in forward citation counts by gender using simple regression analysis to account for non-quality-related patent characteristics. Overall, our evidence points to the undercitation of female lead-authored patents relative to the citations that would have been received by the same patent had its lead author been male.

3.1 Comparing 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 3, and visually demonstrates that females receive fewer citations than males.[11]

On average, male lead inventors receive significantly more citations than female lead inventors (22.6 citations for males, 20 for females, F-stat = 239, Table 1). Testing the difference in distributions, we find a Kolmogorov-Smirnov statistic of $D = 0.034926$, with a p-value of $= 2.2 \times 10^{-16}$, further suggesting that male and female forward citation counts come from different distributions.

We more formally consider the contribution of gender on patent citations by estimat-

[11] Importantly, we note that the number of observations changes between each table as we first balance between male and female and use the propensity network to eliminate outliers.

ing the following OLS model:

$$Y_i = \beta_1 I\left(FemaleInventor_i\right) + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i, \tag{5}$$

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 errors in this paper, unless otherwise noted, are double-clustered by patent issue year and attorney. β_1 is our coefficient of interest, where a positive value would indicate that women receive more citations than males.

The estimates in Panel A of Table 2 present results for the full sample of patents (extensive margin), while Panel B presents the estimates for those patents which receive at least one forward citation (intensive margin).[12] The estimates suggest that female lead investors receive between 0.7 to 2.5 fewer citations than males, depending on specifications and 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 authored patents, the patterns from this simple analysis raise questions. Either the female authored patents being approved are of lower quality, on average than those of males, or, despite the higher bar for approval of female-authored patents, the quality of these patents is not being appropriately reflected in citation counts, with women experiencing undercitation of their inventions relative to males.

3.2 Using C-TEXT to Compute the Treatment Effect

To explore this, we turn to our C-TEXT model. As a reminder, C-TEXT first trains two mappings, one using only patents from male inventors and a second for female inventors. Armed with our two mappings, we pass the male patents through the female mapping, and vice versa. From this, we can estimate the counterfactual number of

[12]Most patents do not receive any forward citations; in general, female lead-authored patents appear to be less likely to receive any citations than those with a male lead author. Panel A of Figure IC2 in the Appendix presents a histogram of the natural logarithm of citations for the intensive margin sample.

citations a patent would have received had its lead author been of the opposite gender, $Forward\widehat{Citation}_i$. We plot the histogram of predicted citations from C-TEXT by gender in Panel B of Figure 3.[13]

The C-TEXT methodology allows us to compute the ATE for treatment with a female lead author, mediating for patent quality through the textual content of the patent. Using our neural network mappings, we calculate the average treatment effect (ATE) of being a female lead author inventor, as defined by (Equation 3), to be -2.66, and the C-TEXT implied ATT (Equation 4) to be -2.38. This means that, on average, having the first author on a patent be female is associated with over 2.5 fewer forward citations compared to having a male lead author for the same patent.

Whiel the ATE and ATT statistics provide us with the overall treatment effect, controlling for the quality/content 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 C-TEXT embeddings would not pick up, such as the art unit or technological classification or the tendencies of the specific patent examiner assigned to the patent. Thus, we next compare the actual forward citations to the model-implied citations in the patent level's absence of treatment (male neural network) in a regression framework that allows us to control for such characteristics.

Specifically, we calculate the difference in actual versus predicted citations as:

$$Delta_i = ForwardCitation_i - Forward\widehat{Citation}_{i|T_i=0}, \tag{6}$$

where $ForwardCitation_i$ is the actual number of citations to a given patent authored by a given gender and $Forward\widehat{Citation}_i$ is the number of citations implied by the C-TEXT model if the lead inventor had been male.

The *Delta* measure is designed to capture the discrepancies in citation counts of patents, both overcitation and undercitation, with an application that spans both male and female-authored works. The measure holds the reference gender constant (in this

[13]For the intensive margin, we plot the histogram of the natural logarithm of citations in Panel B of Figure IC2.

analysis, male) to establish a uniform reference point for evaluating all patents.[14] A negative *Delta* implies that a given patent has amassed fewer citations than the projection established by the quality-adjusted standard posited by the C-TEXT model for this patent if it had a male lead author. As an example, consider a particular patent that has been cited 12 times in actuality and has a computed *Delta* of -3. This *Delta* implies that when we input the text of this specific patent into a neural network calibrated to *predict* its forward citation count assuming a male lead author, the prediction asserts that it should receive 15 forward citations. This projection rests solely on the content of the patent itself and negates the influence of the lead author's gender. It is imperative to underscore that the *Delta* value is not necessarily negative, as depicted in Figure 4, and encompasses broad applicability across patents lead authored by both male and female authors.

We can then relate *Delta* to various patent attributes unrelated to quality, such as examiner, art unit, etc. To do so, we re-estimate models of the type described in Equation 5, replacing the actual number of forward citations for a patent with *Delta*. The estimates in Panel A of Table 3 present estimates for the full sample of patents, including those with zero citations, while Panel B presents the estimates for the subsample of patents that receive at least one citation (intensive margin). Across both panels, the estimates suggest that patents with female lead inventors are undercited relative to what would be expected had the patent remained exactly the same, except that the lead inventor was instead male.

Interpreting our point estimate for the most restrictive specification for the extensive margin, we find that patents with female lead inventors receive 1.4 fewer citations than would be estiamted if they had had a male first author instead, for the same patent, controlling for other characteristics and mediating for quality using C-TEXT. Interpreting these results relative to the sample mean of the number of forward citations, this is over 10%. As expected, the magnitudes are relatively unchanged when including patent-year, art-unit, examiner, attorney, and assignee fixed effects.[15]

[14]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.

[15]The relative stability in estimates suggests that our analysis does not suffer from a correlated omitted variable, Oster (2019).

The results presented up to this point utilize the entire patent text and the Longformer embedding. A natural concern is that Longformer is not specific to (trained on) scientific writing, and as a result may not fully pick up patent quality for mediation purposes. To address this concern, we repeat our analyses utilizing the SciBERT embedding. Because the SciBERT embedding, like all BERT models, has computational difficulties scaling to longer texts, for this exercise, we utilize patent abstracts rather than full patent text. The estimated ATE using SciBERT, presented in Figure 5, is *larger* than that obtained using the Longformer specification. Table IC1 presents estimates of Equation 5 when using the SciBERT embedding version of C-TEXT for the intensive margin sample. The estimates suggest that women are undercited by between -3.1 to -3.5 citations per patent. This is equivalent to roughly 18% relative to the sample mean–an effect even larger than our baseline results. The advantage of the SciBERT transformer for interpretation of the results is that it is trained to assess the quality of scientific writing such as patents; the drawback is that we can only employ it only on patent abstracts due to computational limitations. Still, the use of a scientific writing-specific embedding, if anything, suggests that the effects are still quite large.

The combined results represent causal evidence suggesting 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.[16]

4 Cross Sectional Heterogeneity

Next, we explore whether these patterns of undercitation are uniform across a variety of dimensions of heterogeneity in patent characteristics. For computational tractability, we focus these tests on the sample of patents that receive at least one forward citation.[17]

[16] In the Online Appendix, we show qualitatively similar patterns using the probabilities of a patent being in the top decile of citations. See Table IC2

[17] This choice was made due to computational complexity and taking into account the similarity between both the intensive and extensive estimates in the prior section.

4.1 Patent CPC Section

First, a reasonable question is whether the underciting of female lead inventor patents uncovered in our main models holds across all technology categories or whether there is variation across fields. We next explore this heterogeneity. Specifically, we estimate the following model:

$$\begin{aligned} Y_i &= \beta_1 I\left(FemaleInventor_i\right) + \beta_2 I\left(FemaleInventor_i\right) \times \left(CPCSection\right) \\ &\quad + \delta_{CPCSection} + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i, \end{aligned} \tag{7}$$

where the subscript and notation match the prior estimating equations. As in the main analysis, standard errors are double clustered by year and attorney.

First, we interact our female indicators with the seven CPC sections to study differences by broad field categories. The estimates in Table 4 highlight important heterogeneity across patent CPC sections. Column (1) of Table 4 presents the estimates for the model using the major sections, where the outcome variable is the actual number of forward citations received by the patent. Column (2) re-estimates the model using the difference between the actual citations and the number predicted by the C-TEXT model for the same exact patent when authored by a male, *Delta*.

To interpret the overall effects of our variable of interest for each section, we need to add the coefficients of the indicator for female lead inventors with the interaction term for each major section. In general, we observe some level of undercitation for all patent sections, with particularly large disparities for the Human Necessities section. Put differently, if a female lead-inventor patent instead had a male lead inventor, it would have received significantly more citations, regardless of the technology section, echoing our baseline results.

As an alternative approach, we break down the technology grouping using the NBER subcategories. Specifically, we estimate the following.

$$\begin{aligned} Y_i &= \beta_1 I\left(FemaleInventor_i\right) + \beta_2 I\left(FemaleInventor_i\right) \times \left(Subcategory\right) \\ &\quad + \delta_{PatentSubcategory} + \delta_{GrantYear} + \delta_{Customer \times Examiner} + \varepsilon_i, \end{aligned} \tag{8}$$

To ease interpretation, Figure 6 presents the linear combination of the female lead indicator and the interaction coefficients (Female Lead Inventor $\times$ Subcategory) 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 6, for a large portion of the technology subcategories, the estimates suggest that patents with lead female inventors are cited significantly less than 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 clubby, with higher barriers of entry for female inventors. On the other hand, the underlying forces that lead to undercitation for patents with female first authors 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.

To explore these issues further, we denote a category as an "emerging field" if the art unit first appeared within three years of the patent being granted. We then re-run our models, adding an indicator for an emerging field as well as an interaction between that indicator and the indicator for a female lead inventor. Our coefficient of interest is the interaction between the indicator for female inventors and emerging fields.

The estimates are presented in Table 5. Panel A presents estimates where the LHS variable is actual forward citations, and Panel B presents estimates where the LHS 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. However, when moving to Panel B, our estimates show that the main result still holds — patents with female first authors exhibiting an estimated 1.3 to 1.7 fewer citations than would be predicted if the first author had been male, depending on specification. That said, the estimates also suggest that women receive even fewer citations in emerging fields relative to their male counterparts. Here, we see that they receive 2.1 to 2.7 fewer citations, suggesting that new fields exhibit the same general pattern of undercitation for female inventors, rather than new fields reducing barriers or bias.

4.3 Time Since Patent Grant

A natural question is whether the undercitation we observe above is present from the outset or whether it primarily materializes or diminishes later in the life of the patent. On the one hand, undercitation may be present from the outset but diminishes over time as inventors and examiners become more familiar with the patent and its quality. Alternatively, the bias may increase and become more pronounced over time, potentially indicating a self-reinforcing effect that could be harder to overcome. Examining the timing of the bias in citations can provide valuable insights into the nature of the undercitation of female inventors and inform potential interventions to address this issue.

To investigate the timing of undercitation, we create separate samples of forward citations based on the number of years that have passed since a given patent was granted. Specifically, we divide the post-grant period into four sub-periods: [0-1) years post-grant, [1-5) years, [5-10) years, and [10-20] years. For each of these sub-periods, for each patent, we collect the forward citations the patent receives during this sub-period. For each sub-period, we then re-run our analysis and calculate the *Delta* in citations (actual minus predicted by the male model) after mediating for patent quality.

The estimates are presented in Table 6. Column (1) presents estimates from forward citations to patents received in the first year after the patent grant, column (2) presents estimates for forward citations received in years 2 to 5 after patent grant, column (3) for citations received in years 6 to 10, and column (4) years 11-20. In each column, the dependent variable is the *Delta* estimated from C-TEXT using only forward citations received during that subperiod (by necessity, the number of observations is smaller in later sub-periods as fewer of the patents in our sample will yet have histories of that length. Also, few patents receive citations in their first year, resulting in a smaller sample). As can be seen from the estimates in the table, the undercitation for patents with a female first author relative to what would be expected if the first author had been male increases over time since the patent grant, consistent with undercitation being self-reinforcing over time. The coefficient for lead female inventors is economically and statistically insignificant in the first period ([0-1) years), but becomes more pronounced and strongly statistically significant in the subsequent periods, with estimates of -0.65, -0.60, and -2.0 for the [1-5) years, [5-10) years, and [10-20] years periods, respectively. These findings are consistent with the notion that undercitation in later periods may be further reinforced by prior undercitation, leading to a situation in which overcited patents continue to be overcited and the *Delta* becomes larger over time.

4.4 Evolution of Undercitation Over the Sample Period

The estimates we present in the prior analyses suggest that across fields, patents with female first authors are consistently undercited relative to what would be expected for the same patent had its first author been male. A natural question is whether these patterns vary over time, as gender norms, female participation in the workforce, and in academia has changed over time (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. Here, to ensure we are comparing apples to apples, 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 is larger later in the life of the patent, as we showed above, 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 must exclude 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. Next, we apply our C-TEXT methodology to calculate *Delta* for each patent. Importantly, for this test, we train our model using forward citation counts for a ten-year period after the patent was granted to avoid biases in our calculations.

We then estimate models of the following nature:

$$\begin{aligned} Y_i \;=\; & \beta_1 I\left(FemaleInventor_i\right) \\ & + \sum_{j=1977}^{2011} \beta_j I\left(FemaleInventor_i\right) \times I\left(GrantYear = j\right) \\ & + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i, \end{aligned} \tag{9}$$

where β_1 estimates the average undercitation of females across the entire sample, and the set of coefficients β_j estimate the marginal *Delta* between actual forward citations and the number of forward citations that would be predicted had the first author of the patent has been male in each patent grant year, with 1976 as the year of comparison. The omitted group is 1976.

To ease interpretation, Figure 7 presents the linear combination of the female lead indicator and the interaction coefficients (Female Lead Inventor × Grant Year) graphically.[18] From the figure, we observe clearly that the average undercitation of patents with female lead authors has been persistent over time.

[18] The time-invariant estimate for the coefficient on the female lead inventor variable is -1.18.

5 Who Undercites Female Inventors?

So far, we have presented causal evidence that patents with female lead inventors receive fewer citations than the same patents would be estimated to receive with male lead inventors. Next, we explore the source of the undercitation: whether it is driven by inventors or examiners, and the role of their gender.

To set the stage for this analysis, we first discuss how a citation is added to a patent. When applying for a patent, applicants cite supporting patents whose inventions the current patent is building on top of. If, however, the patent examiner deems that there are additional relevant citations that have not been included by the inventor, the examiner will add these to the patent application. As a result, the documented undercitation of patents with female lead inventors may stem from the original inventor-added citations, additional examiner-added citations, or a combination of both.

To explore the source of the undercitation, 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 then decompose the sources of undercitation of female lead-inventor patents.

We begin our analysis by studying examiner-added citations. For a given patent, we take all forward citations that occur due to being added to a future patent application by an examiner. We then break these into forward citations added by female examiners and forward citations added by male examiners. Following similar logic to our main tests, we then apply the C-TEXT model, estimating a neural net for male-lead inventor patents and a neural net for female-lead inventor patents to estimate forward citation counts added by examiners of each gender based on the gender of the lead inventor on the patent of interest. We then employ the C-TEXT methodology to mediate for the quality of the patent and calculate the *Delta* between actual forward citations and what would have been estimated by the examiner neural net for the same patent if it had a male lead author.

Table 7 presents the results of the estimation of regression models, using the C-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 suggest that female examiners, when adding citations to patents, do not appear to *undercite* female lead inventor patents. Put differently, we cannot reject the hypothesis that patents with a lead female inventor receive similar citation "add" from female examiners as they would be expected to had their first author instead been male. The coefficient estimates are economically small, ranging from -0.06 to 0.01 citations, with no statistical significance. Panel B repeats the analysis for male examiners. Altogether, the estimates suggest that patents with a lead female inventor also receive similar citation "adds" from male examiners as they would be expected to had their first author instead been male. Here, some coefficients attain statistical significance, but the coefficients are small, and economic magnitudes are negligible. The coefficient estimates range from -0.075 citations to -0.03, with the economic magnitudes being less than 2% of the sample mean. Overall, the estimates suggest minimal, if any, undercitation of female lead inventor patents by male examiners. Taken together, the estimates suggest that the undercitation we observe for female lead inventor patents is not driven primarily by examiner patents.

Having established this fact, 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 a patent by male and female inventors and that which would be predicted for the same patents if their lead inventor had instead been male. The estimates from the C-TEXT *Delta* regressions are presented in Table 8. 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. First, as can be seen in Panel A, undercitation of female lead inventor patents does not appear to be due to other female lead inventors citing such patents less than they would have had the patents have been male authored. The estimates in columns (1) through (4) are statistically insignificantly different from zero and of minuscule economic magnitude. In contrast, however, in Panel B, we observe a clear pattern of statistically significant negative coefficients of interest. Put differently, male inventors are significantly less likely to include a citation to a female lead authored patent than they would be predicted to have that same exact patent had a male lead inventor. The estimates in Panel B suggest undercitation of female lead inventor patents by male inventors by 1.4 citations, or 11% of the sample mean. This is large, both economically and statistically, and comparable in overall magnitude to the overall effect of -1.76 citations documented in Table 3.

To further study this pattern we present the ATE for each of the four subgroups. Presenting the treatment effects in Figure 8 we again see a similar pattern. That is, the ATE for male inventors the largest of the subgroups, estimates to be -1.73. These are followed by female lead inventors with an estimate of -0.52. Again, we see that examiners tend to be more even-handed, with female and male lead inventors having an ATE of -0.17 and -0.06.

Taken together, the results strongly suggest that the undercitation of female lead patents (relative to what would be expected had their lead inventor been a male instead) is primarily driven by male lead inventors underciting patents with female lead inventors in their patent applications. Of course, such undercitation does not necessarily imply discrimination on the part of male inventors. Alternative explanations could include men having more and stronger connections to, or familiarity with, other male inventors and, as a result, being more familiar with patents filed by other male lead inventors. Future research will be necessary to fully distinguish the reason for the undercitation.

6 Robustness

6.1 Definition of "Female Authored Patent"

A potential concern is that the manner in which we classify patents by gender using solely the lead inventor on the patent in some unknown way produces spurious results. Understanding the effect of different classifications for a "female" patent is useful in shedding light on whether in fact our estimates are due to a gender treatment effect, or some other mechanism. In the analysis up to now, we use the name of the first inventor to assign author 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, and not just the first author, when attributing the gender of authors to the patent.

Importantly, our estimates are robust to a number of alternative approaches to attributing author gender to a patent. First, we limit our sample to patents with only one author, comparing female sole-authored patents to male sole-authored patents. This shuts down concerns that the gender of additional authors other than the lead author may be driving the baseline results.[19] The single-author sample also exhibits consistent undercitation of female inventor patents, as can be seen from Panel A of Table IC3. The estimates suggest that women are undercited by over 2 patents per citation in this sample, nearly 11% of the sample mean. These results are highly significant and in line with the baseline results.

Second, we construct a new sample that includes only patents with a majority of inventors of the same gender, both single authors and teams. Here, again, we find similar patterns of undercitation of female authored patents, as can be seen in Panel B of Table IC3. Estimates from this test suggest that patents with a majority of female inventors receive 0.4 to 0.8 fewer citations than they would have had they instead been a majority male inventor. Although statistically significant, the economic magnitudes are smaller than those estimated in our main models. The smaller coefficient sizes would be

[19]Importantly, the neural network that measures the propensity for a given patent to have been written by a male versus a female assures that what we are picking up in the C-TEXT *Delta* is quality as indicated by text content, as opposed to writing style.

consistent with a world in which the gender of the lead inventor is particularly salient.

6.2 Placebo Test

To further strengthen our conclusions, we further 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" authored and the other to be "female" authored. We then apply the C-TEXT analytical procedure to ascertain the role, if any, that may be played by the *methodology* in shaping our results, rather than the data. As expected, Table IC4 shows clearly that, regardless of the model specification, there appears to be no statistical correlation between the assignment to a "female" patent and the resulting *Delta*. All in all, the placebo tests suggest that the C-TEXT methodology itself is not likely to be driving our results.

6.3 Training Period

A second concern is that some of the patents in the earliest part of the sample may have been particularly important or influential and that this happens in a period where perhaps females are patenting less, and this in some way influences the results, producing spurious estimates of undercitation. To address this concern, we split the sample into two sub-parts: one for patents issued pre-2000, and another set for those issued post-2000.

We then re-estimate our baseline specification on each of the subsamples and present our results in Table IC5. Panel A, which presents estimates for the pre-period, suggests an average effect of -1.82, similar to our baseline effect estimated in Table 3. Similarly, in the post-2000 sample estimates, presented in Panel B, we again find similar estimates, with a slightly smaller but still significant main effect of -1.16 citations for a female lead authored patent relative to what it would be predicted to have received had the lead inventor instead had a male name. Thus, it appears that our estimates are not a spurious result of our choice of the longer sample period.

6.4 Overfitting of Model

A standard concern with these types of models is overfitting the training data. In our setting, we train two different models by completing multiple passes of our training dataset through our algorithm. Each pass of the data in these settings is referred to as an epoch. We estimate our neural networks using 20 epochs. While numerous epochs help improve the predictive probability of the neural networks, they run the risk of overfitting our model to the data. Such overfitting would then result in relatively poor out-of-sample performance. In the context of our paper, this would result in incorrect or biased out-of-sample estimates of the number of citations.

We address this concern by studying the loss function, as presented in Figure IC3, 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 two key pieces of evidence that suggest we have not overfit the model. First, as we increase the number of epochs, the MSE tends to decrease, indicating that each additional pass adds information to the estimation. Second, we find diminishing improvements to the error rate as we approach 20 epochs, the number of passes utilized in our estimations. Taken together, these findings suggest that our model is unlikely to be overfitted and, as a result, that reasonable counterfactual citations are estimated from our neural networks.

7 Economic Value of Patent and Citation Bias

We next consider the relationship between the economic importance of a patent, as evaluated by public markets at the time of issue, and forward citations. As we have previously discussed, undercitation of female-authored patents tends to persist over time. In contrast, the expected economic value of a patent, as assessed by public markets, is forward-looking and can be determined at the time of issuance. An interesting question is then whether these forward looking market estimates of a patent's economic value relate more closely to actual forward citations, or to the estimated number of forward citations we obtain out of C-TEXT.

To test this, we employ the subsample of female lead-authored patents and employ the patent-level measure of economic value proposed in Kogan et al. (2017). This measure is computed for patent issues for publicly-traded U.S. firms and utilizes the stock market's response to news about patent grants. Specifically, we use the log value of innovation, deflated to 1982 (million) dollars using the CPI. We then examine whether this measure, when regressed on actual forward citations and the predicted counter-factual forward citations if the first author had been male, loads on one of these measures in particular versus the other. Specifically, we estimate:

$$\begin{aligned} log(Dollar_i) &= \beta_1 ForwardCitation_i + \beta_2 Forward\widehat{Citation}_{i|T_i=0} \\ &\quad + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Examiner} + \delta_{Attorney} + \delta_{Assignee} + \varepsilon_i \end{aligned} \tag{10}$$

The estimates are presented in Panel A of Table 9. Across all specifications,the coefficient on $Forward\widehat{Citation}_i$ loads significantly, while the coefficient on actual forward citations does not have statistical significance. The estimates suggest that the market does not appear to undervalue female-authored patents relative to what it would have had that same patent been authored by a male lead inventor. The estimates provide support to the notion that actual forward citations for female lead authored are biased downwards due to the gender of the lead inventor on the patent.

8 Discussion and Conclusion

We provide causal 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 female-authored patents, relative to what that same patent would have received had the first author 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 female authored patents are undercited, and as a result, female compensation 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 et al., 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 C-TEXT methodology for causal inference. 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. Causal inference using text 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.

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, 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)* .

Cook, Lisa, 2020, Policies to broaden participation in the innovation process, *Policy Proposal, The Hamilton Project, Brookings Institution, Washington, DC* .

Cook, Lisa D, 2014, Violence and economic activity: evidence from african american patents, 1870–1940, *Journal of Economic Growth* 19, 221–257.

Cook, Lisa D, and Chaleampong Kongcharoen, 2010, The idea gap in pink and black, Technical report, National Bureau of Economic Research.

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

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, Jean-Philippe Garant, Hannah Herman, and David J Munroe, 2013, Why are women underrepresented amongst patentees?, *Research Policy* 42, 831–843.

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, 2022, Does finance benefit society? a language embedding approach, Working paper.

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

Koffi, Marlène, 2021b, 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.

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

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.

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, Kunal Sachdeva, and Aaron Yoon, 2022, The evolution of esg reports and the role of voluntary standards, Technical report.

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

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.

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.

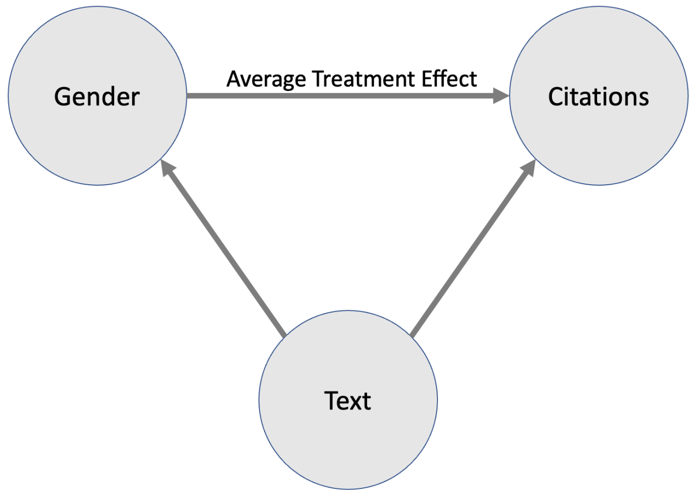

(a) Panel A: Model for ATT, Confounding

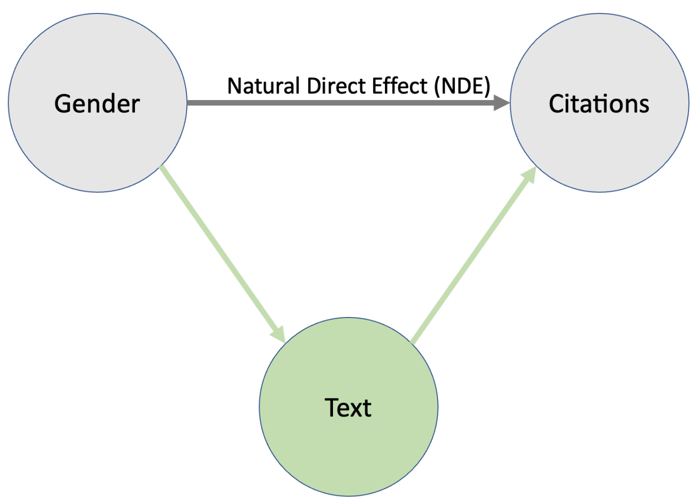

(b) Panel B: Model for NDE, Mediating

FIGURE 1: TEXT, GENDER, CITATIONS

Citation is the outcome of interest, *Gender* is the treatment, and *Text* are the sequence of words. Panel A depicts the average treatment effect, with the assumption that *Text* carries sufficient information to adjust for confounding (common cause) between outcome and treatment. Panel B depicts the natural direct effect (NDE), where the text is a mediator of the treatment on outcome.

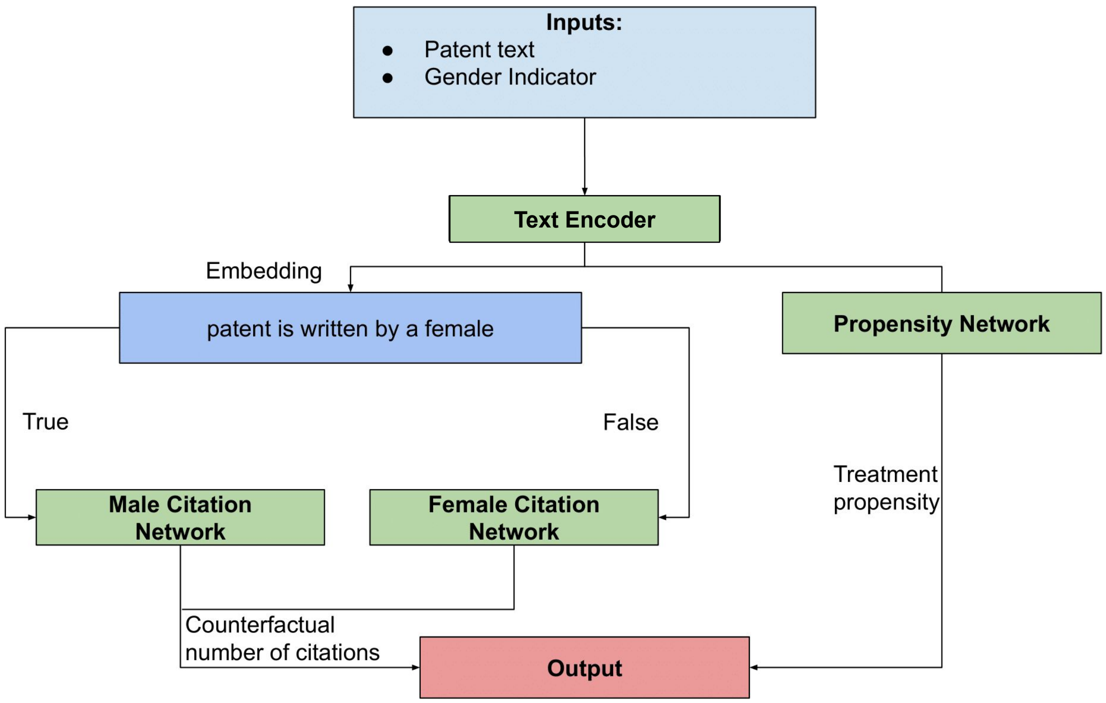

FIGURE 2: C-TEXT ESTIMATION PROCEDURE

The figure illustrates the estimation procedure of C-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 counterfactual estimation. Finally, the red block is the output that combines the outputs of the citation estimation networks and the propensity score estimation network.

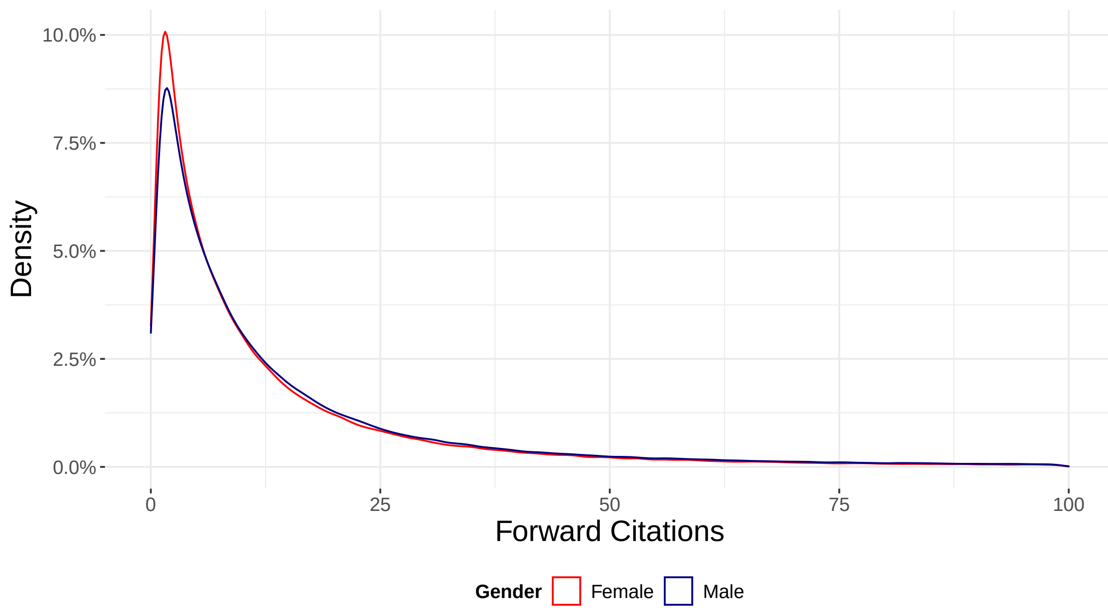

(a) Panel A: Observed Forward Citations

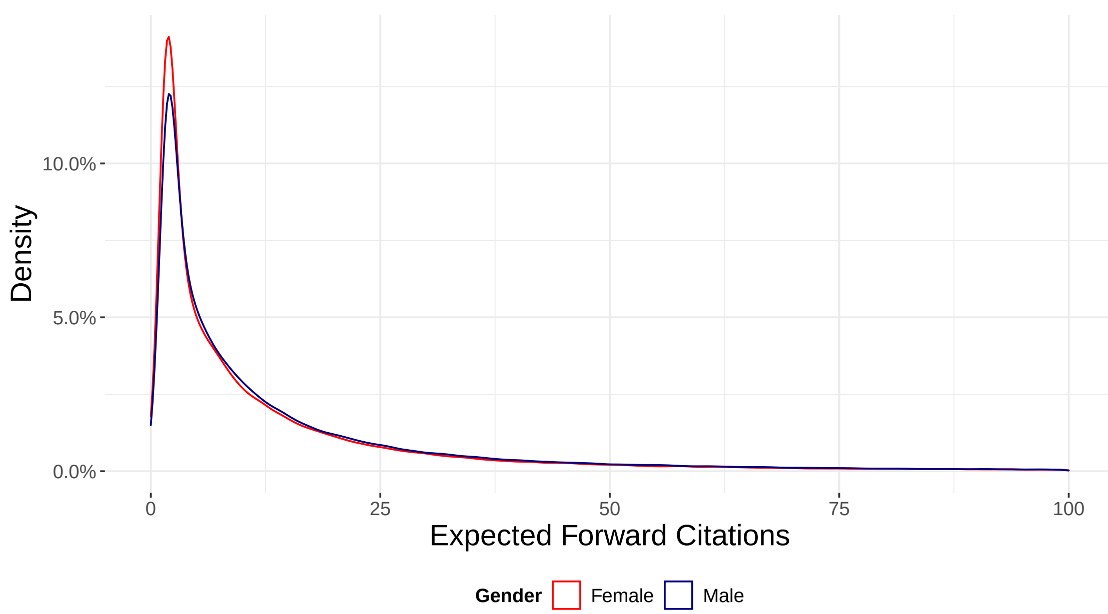

(b) Panel B: C-TEXT Implied Forward Citations

FIGURE 3: 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 C-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 IC2.

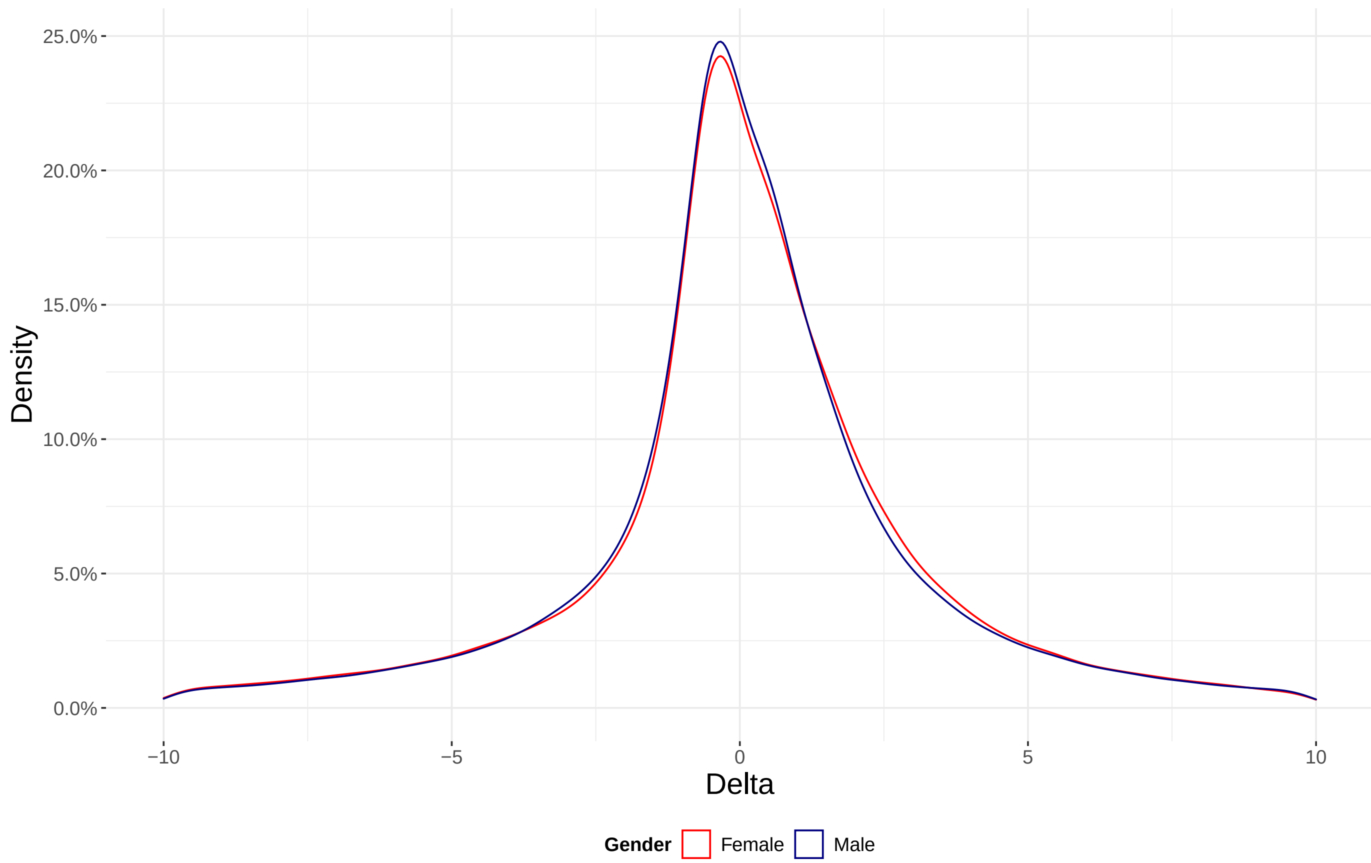

FIGURE 4: DISTRIBUTION OF DELTA

This figure illustrates the difference between forward citations and expected forward citations for male lead authored patents, as defined by Equation 6. The red line corresponds to females, the blue line corresponds to males. The distribution is truncated between -10 to 10.

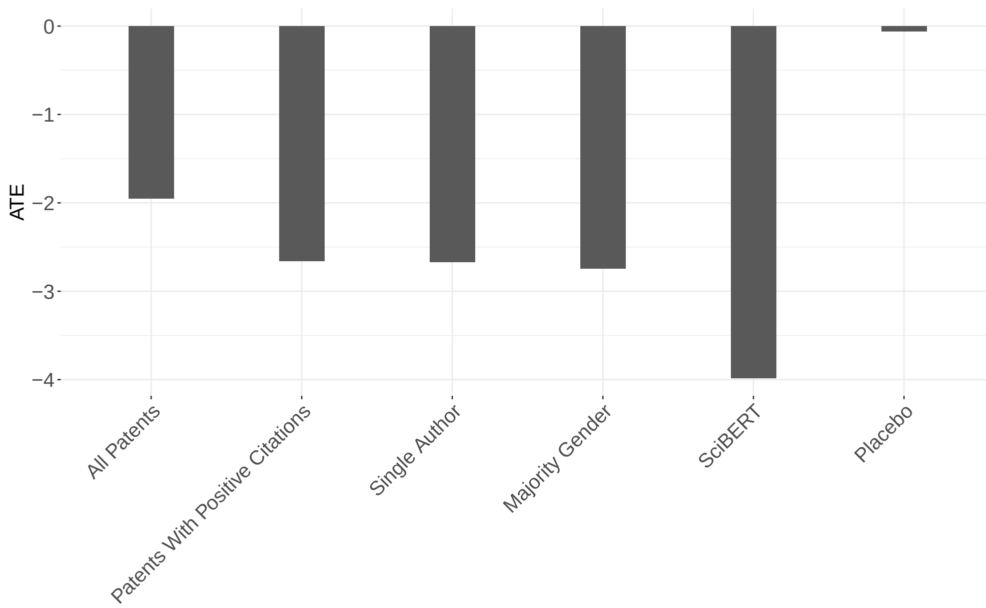

FIGURE 5: AVERAGE TREATMENT EFFECT FOR DIFFERENT SAMPLES

This figure plots the average treatment effect (ATE), as defined by Equation 3. Each bar represents a different sample and estimate. The first bar corresponds to all patents. The second bar corresponds to patents with a positive number of citations. The third bar corresponds to patents with a single author. The fourth bar corresponds to teams with a majority of a given gender. The fifth bar uses applies SciBERT on the abstracts of patent text. The sixth bar is the placebo test.

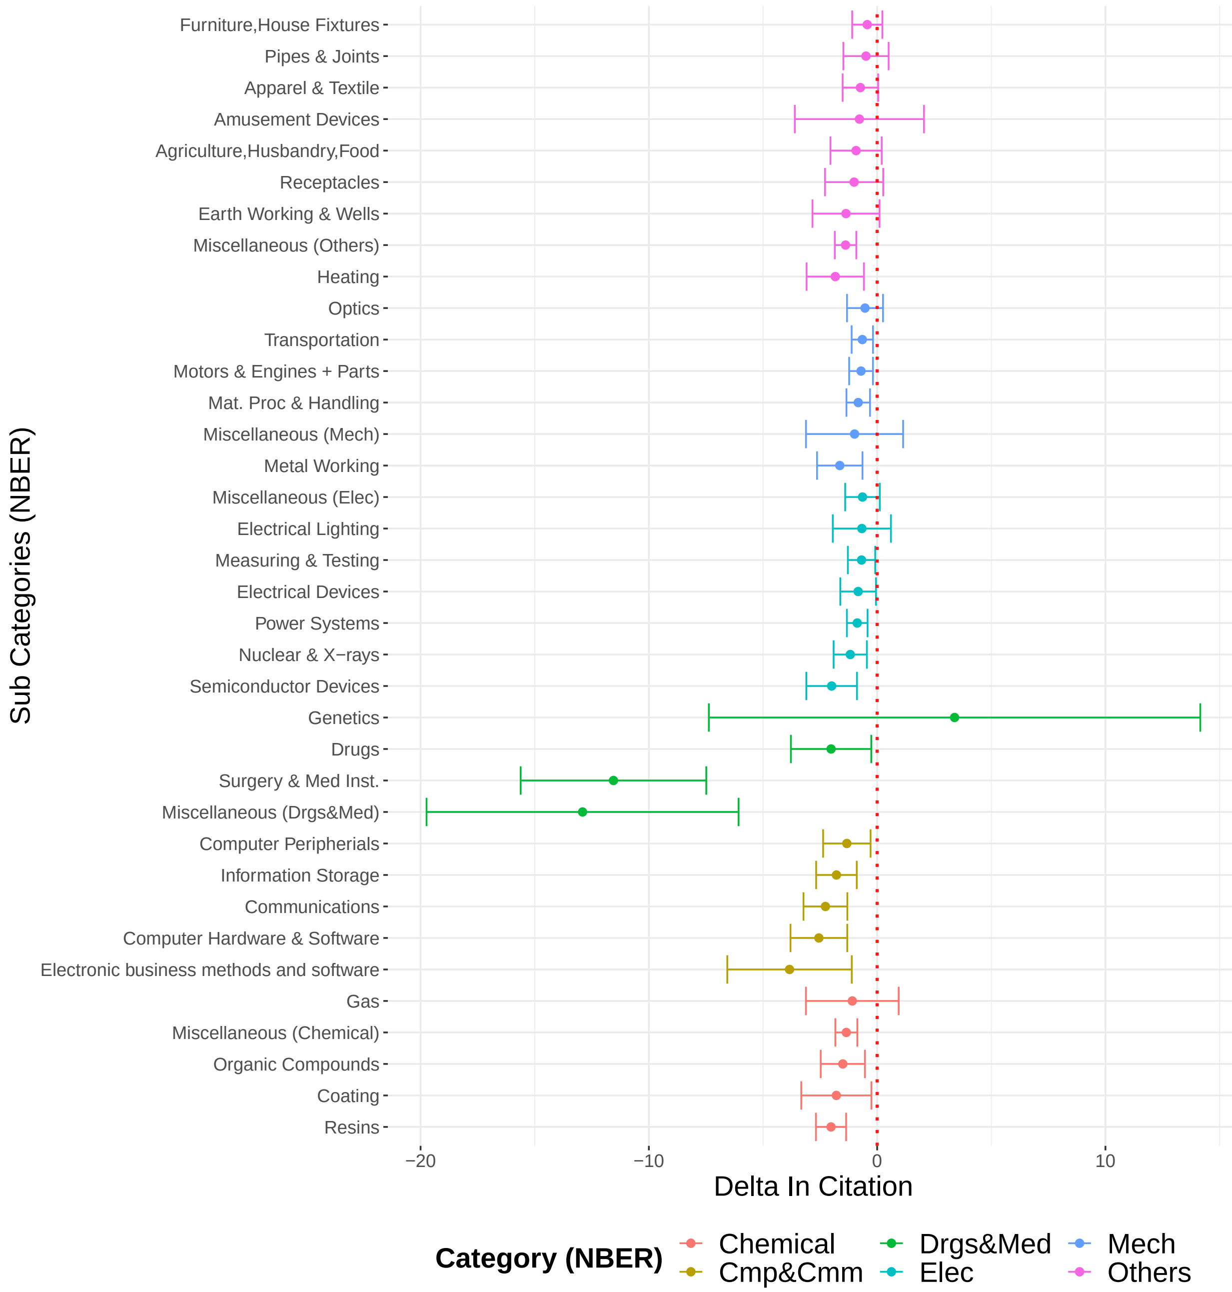

FIGURE 6: DELTA IN CITATIONS BY PATENT SUB CATEGORIES

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.

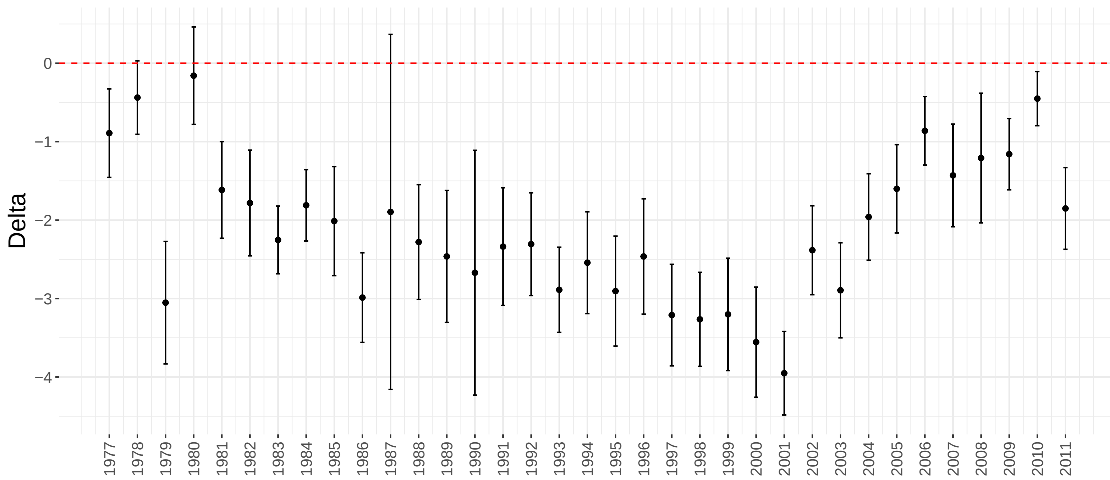

FIGURE 7: EVOLUTION OF DELTA OVER TIME

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. 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.

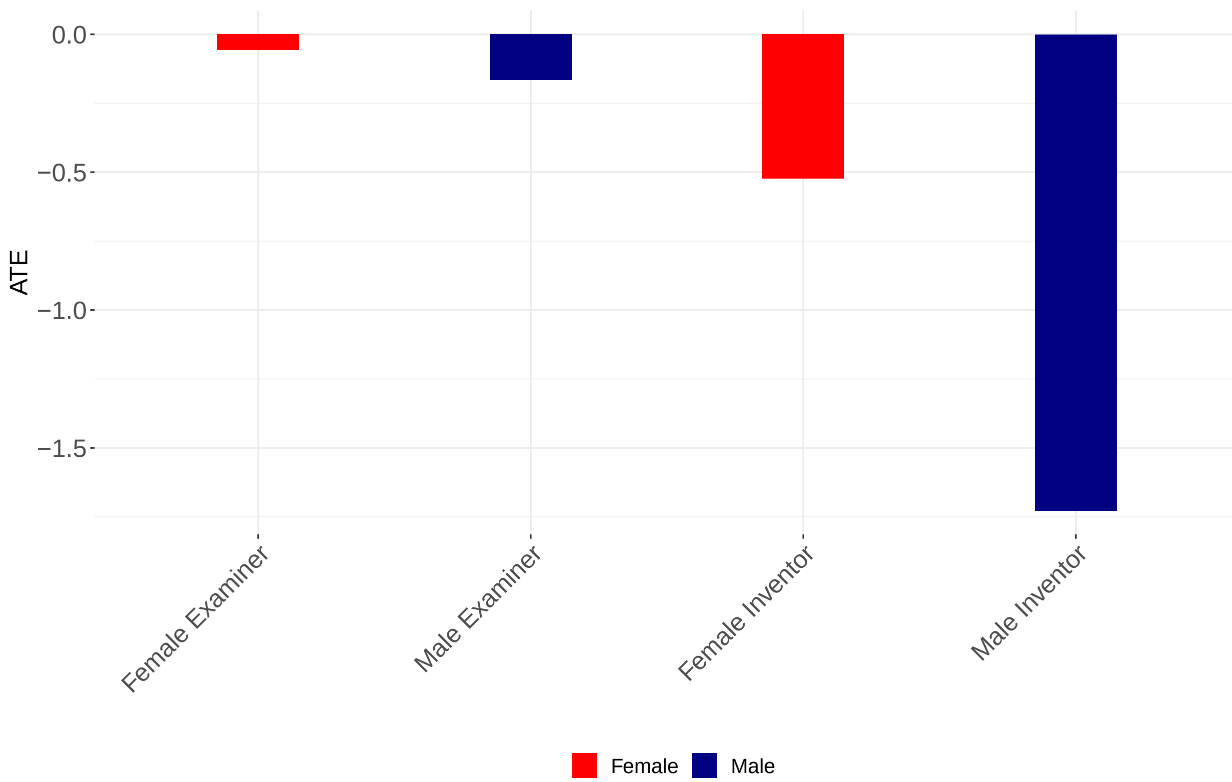

FIGURE 8: AVERAGE TREATMENT EFFECT, EXAMINER AND INVENTOR ADDED CITATIONS

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.

TABLE 1: SUMMARY STATISTICS

This table provides summary statistics on patents and citations. The sample covers patents issued from 1976-01-01 through 2021-12-31. 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%, 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	294004	22.58	60.55	240993	20.07	57.23	F= 239.299***
Panel B: Difference By Top Decile Innovations							
Breakthrough	294004			240993			$\chi^2 = 421.368^{***}$
→ No	262513	89%		219250	91%		
→ Yes	31491	11%		21743	9%		
Panel C: Difference by Cooperative Patent Classification							
CPC Section	293965			240955			$\chi^2 = 2268.912^{***}$
→ Chemistry	30171	10%		29724	12%		
→ Electricity	67709	23%		61192	25%		
→ Fixed Constructions	9315	3%		5783	2%		
→ Human Necessities	37546	13%		32667	14%		
→ Mechanical Engineering	24915	8%		15649	6%		
→ Performing Operations	48795	17%		34626	14%		
→ Physics	72321	25%		58868	24%		
→ Textiles	3193	1%		2446	1%		

TABLE 2: LEAD FEMALE INVENTORS, NO C-TEXT ADJUSTMENT

This table reports estimates of Equation 5 and studies the number of forward citations by the gender of the lead inventor. The dependent variable is the forward citation for each patent. Panel A uses the sample of all patents while Panel B uses patents with a positive number of forward citations. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Extensive Margin

	Forward Citations					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-3.248***	-1.695***	-1.019***	-1.056***	-0.781***	-0.2817**
	(0.169)	(0.135)	(0.122)	(0.117)	(0.114)	(0.0967)
Intercept	15.70***					
	(0.48)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	22.9%	11.9%	7.2%	7.4%	5.5%	2%
Observations	1039516	1039516	1039516	1039516	1039516	1039516
R^2	0.001	0.096	0.125	0.179	0.266	0.451

Panel B: Intensive Margin

	Forward Citation					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-2.511***	-2.018***	-1.87***	-1.833***	-1.373***	-0.722**
	(0.405)	(0.334)	(0.35)	(0.345)	(0.324)	(0.253)
Intercept	22.58***					
	(1.47)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	11.7%	9.4%	8.7%	8.5%	6.4%	3.4%
Observations	534997	534997	534997	534997	534997	534997
R^2	0.000	0.088	0.099	0.157	0.256	0.460

TABLE 3: LEAD FEMALE INVENTORS, C-TEXT (LONGFORMER)

This table reports estimates of Equation 5 and studies the number of forward citations by the gender of the lead inventor. 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 6. Panel A uses the sample of all patents while Panel B uses patents with a positive number of forward citations. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Extensive Margin

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-1.0937***	-1.2513***	-1.3569***	-1.3677***	-1.4286***	-1.436***
	(0.0912)	(0.0977)	(0.0926)	(0.0937)	(0.0891)	(0.105)
Intercept	-1.4126***					
	(0.0783)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	7.7%	8.8%	9.6%	9.6%	10.1%	10.1%
Observations	1039516	1039516	1039516	1039516	1039516	1039516
R^2	0.001	0.017	0.023	0.055	0.127	0.257

Panel B: Intensive Margin

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-1.316***	-1.504***	-1.521***	-1.61***	-1.729***	-1.759***
	(0.243)	(0.236)	(0.237)	(0.24)	(0.219)	(0.207)
Intercept	-1.265***					
	(0.135)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	6.1%	7%	7.1%	7.5%	8.1%	8.2%
Observations	534997	534997	534997	534997	534997	534997
R^2	0.001	0.025	0.025	0.068	0.154	0.305

TABLE 4: CITATION BY CPC SECTION

This table estimates the difference in citations by CPC Section. Column (1) uses the number of forward citations as its dependent variable, while Column (2) uses *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 6. Estimates include interactions for the patent category based on CPC Section. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Dependent variable:	
	Forward Citations	*Delta*
	(1)	(2)
Lead Female Inventor	-0.678	-1.278***
	(0.456)	(0.278)
Electricity × Lead Female Inventor	0.267	-0.123
	(0.540)	(0.347)
Fixed Constructions × Lead Female Inventor	-0.881	0.269
	(0.888)	(0.414)
Human Necessities × Lead Female Inventor	-2.17*	-3.266***
	(0.83)	(0.736)
Mechanical Engineering × Lead Female Inventor	0.856	0.496
	(0.525)	(0.339)
Performing Operations × Lead Female Inventor	0.643	0.217
	(0.514)	(0.272)
Physics × Lead Female Inventor	-0.102	-0.546+
	(0.496)	(0.306)
Textiles × Lead Female Inventor	-2.33+	0.669
	(1.36)	(0.475)
CPC Section FE	Yes	Yes
Art Group FE	Yes	Yes
Patent Issue Year FE	Yes	Yes
Examiner FE	Yes	Yes
Attorney FE	Yes	Yes
Assignee FE	Yes	Yes
Observations	534920	534920
R^2	0.460	0.305

TABLE 5: 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 three years of the patent being granted. Panel A uses the forward citation for each patent as its dependent variable. Panel B uses C-TEXT with Longformer, with the dependent variable being 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 6. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: No C-TEXT Adjustment

	Forward Citation					
	(1)	(2)	(3)	(4)	(5)	(6)
Emerging Field	13.16***	14.22***	6.18***	4.24**	3.94**	2.48*
	(3.54)	(1.73)	(1.23)	(1.34)	(1.23)	(1.13)
Lead Female Inventor	-2.605***	-2.120***	-1.972***	-1.933***	-1.486***	-0.796**
	(0.415)	(0.336)	(0.355)	(0.348)	(0.326)	(0.253)
Emerging Field × Lead Female Inventor	0.836	1.09	0.816	1.06	2.46	2.09
	(1.717)	(1.43)	(1.454)	(1.61)	(1.75)	(1.53)
Intercept	22.57***					
	(1.48)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	3.9%	5%	3.8%	4.9%	11.3%	9.6%
Observations	514675	514675	514675	514675	514675	514675
R^2	0.002	0.089	0.099	0.158	0.257	0.462

Panel B: C-TEXT (LONGFORMER)

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Emerging Field	-2.17***	-1.072*	-0.286	0.581	0.714	0.972
	(0.58)	(0.531)	(0.556)	(0.664)	(0.657)	(0.650)
Lead Female Inventor	-1.302***	-1.489***	-1.506***	-1.589***	-1.710***	-1.73***
	(0.243)	(0.234)	(0.235)	(0.236)	(0.212)	(0.20)
Emerging Field × Lead Female Inventor	-2.135**	-2.210**	-2.176**	-2.49**	-2.640**	-2.77**
	(0.752)	(0.772)	(0.776)	(0.85)	(0.951)	(1.00)
Intercept	-1.268***					
	(0.133)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	9.8%	10.2%	10%	11.5%	12.2%	12.8%
Observations	514675	514675	514675	514675	514675	514675
R^2	0.001	0.025	0.025	0.069	0.154	0.306

TABLE 6: YEARS AFTER PATENT IS GRANTED

This table estimates Equation 5 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 author was male, as defined by Equation 6. 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 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level.***, **, *, + denote significance at the .1%, 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	0.22	-0.647***	-0.602***	-2.00***
	(0.26)	(0.085)	(0.084)	(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
Relative to Sample Mean	11.9%	14.4%	8.8%	10.3%
Observations	14510	286787	265095	213540
R^2	0.637	0.270	0.235	0.516

TABLE 7: EXAMINER-ADDED CITATIONS

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 6. 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 1976-01-01 through 2021-12-31. Note, the source of citations is only available following the start of 2001. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Citation Added by Female Lead Examiner

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	0.0096	0.0066	0.0151	0.0118	-0.0478	-0.0685
	(0.0263)	(0.0277)	(0.0287)	(0.0382)	(0.0572)	(0.0850)
Intercept	0.1674***					
	(0.0288)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	0.5%	0.3%	0.7%	0.6%	2.3%	3.3%
Observations	12584	12584	12584	12584	12584	12584
R^2	0.000	0.169	0.173	0.521	0.791	0.928

Panel B: Citation Added by Male Lead Examiner

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-0.071**	-0.0740**	-0.0609*	-0.0688**	-0.0599*	-0.0254
	(0.023)	(0.0227)	(0.0240)	(0.0248)	(0.0246)	(0.0301)
Intercept	0.2347***					
	(0.0266)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	2%	2%	1.7%	1.9%	1.7%	0.7%
Observations	127719	127719	127719	127719	127719	127719
R^2	0.000	0.014	0.019	0.131	0.284	0.516

TABLE 8: 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 author was male, as defined by Equation 6. 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 1976-01-01 through 2021-12-31. Note, the source of citations is only available following the start of 2001. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Citation Added by Female Lead Inventors

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	0.186	-0.045	0.0245	0.135	0.258	-0.0535
	(0.236)	(0.208)	(0.2021)	(0.367)	(0.446)	(1.9293)
Intercept	0.390*					
	(0.156)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	1.9%	0.5%	0.3%	1.4%	2.7%	0.6%
Observations	7209	7209	7209	7209	7209	7209
R^2	0.000	0.258	0.269	0.671	0.856	0.900

Panel B: Citation Added by Male Lead Inventors

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-1.427***	-1.47***	-1.474***	-1.453***	-1.539***	-1.390***
	(0.227)	(0.22)	(0.219)	(0.224)	(0.199)	(0.194)
Intercept	-0.0732					
	(0.0711)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	11.1%	11.4%	11.4%	11.3%	11.9%	10.8%
Observations	333673	333673	333673	333673	333673	333673
R^2	0.001	0.012	0.013	0.072	0.188	0.306

TABLE 9: FORWARD CITATIONS AND VALUE OF PATENT

This table studies the relationship between the measures of citations and the market-implied value of patents. The dependent variable for both panels use the log value of innovation, deflated to 1982 (million) dollars using the CPI, as calculated in Kogan et al. (2017). The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

	log($dollar$)					
	(1)	(2)	(3)	(4)	(5)	(6)
$log(ForwardCitation)$	-0.084	-0.027	0.0073	0.0075	0.027	0.0111+
	(0.059)	(0.039)	(0.0322)	(0.0333)	(0.016)	(0.0056)
$log(Forward\widehat{Citation})_{i\|T_i=0}$	0.324***	0.257***	0.25***	0.241***	0.080***	0.0012
	(0.042)	(0.032)	(0.03)	(0.031)	(0.017)	(0.0057)
Intercept	0.18					
	(0.18)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Observations	87806	87806	87806	87806	87806	87806
R^2	0.021	0.160	0.179	0.315	0.749	0.932

INTERNET APPENDIX
FOR ONLINE PUBLICATION

Appendix A Explanation of Causal Text Analysis (C-TEXT)

C-TEXT is a neural network-based architecture that estimates counterfactuals of a binary treatment where all of the covariates needed for causal identification are contained within a given text. To use C-TEXT to identify the effect of gender on the impact of patents, we first need to train the model. As shown in Figure IB1, the input data for training contains three types of information: the texts of patents, gender indicators of the author(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 to treatment propensities, and two 2-layer perceptrons that map from embeddings to male and female predicted number of citations, respectively. The final loss function is a weighted average of the losses of these four neural networks. After the model is trained, we can use it to estimate the counterfactual number of citations of male written patents if they were written by females and vice versa. As shown in Figure 2, to estimate these counterfactuals, we run the trained C-TEXT model where the input data contains the texts of the patents and gender indicators of the author(s). The texts are first passed through the encoder model to generate a vector embedding for each patent. Then each embedding-gender pair is passed through a decision step: if the author(s) are male, the embedding is passed to the female citations network and if it is written by female(s), the embedding is passed to the male citation network. The counterfactual number of citations is then computed by these two networks. In parallel, regardless of the gender indicators, each embedding is passed through the propensity network to estimate the treatment propensity of this patent. Finally, the output of the model is a set of counterfactual citation-treatment propensity pairs that each correspond to one patent.

Appendix B Encoder Architecture

The encoder architecture works as follows. Let W denote the original input sentence in words. As shown in Figure IB2, 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 output 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 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 entire texts that represent not only the meaning of the tokens but also the contextual relationship between tokens and sentences.

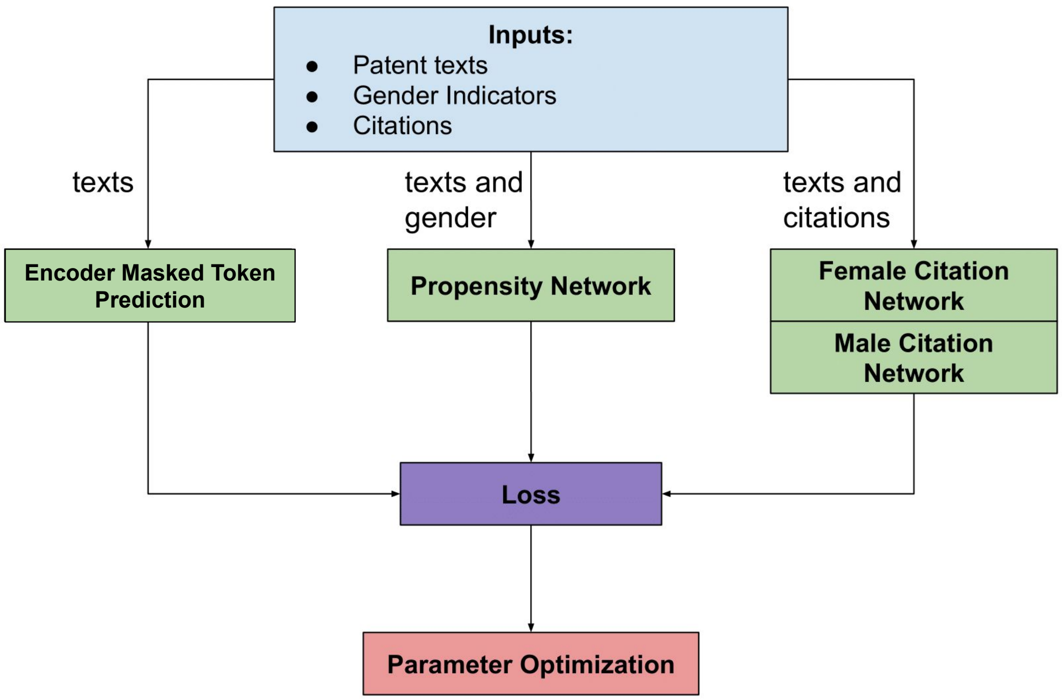

FIGURE IB1: C-TEXT TRAINING PROCEDURE

The figure illustrates the training procedure of C-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.

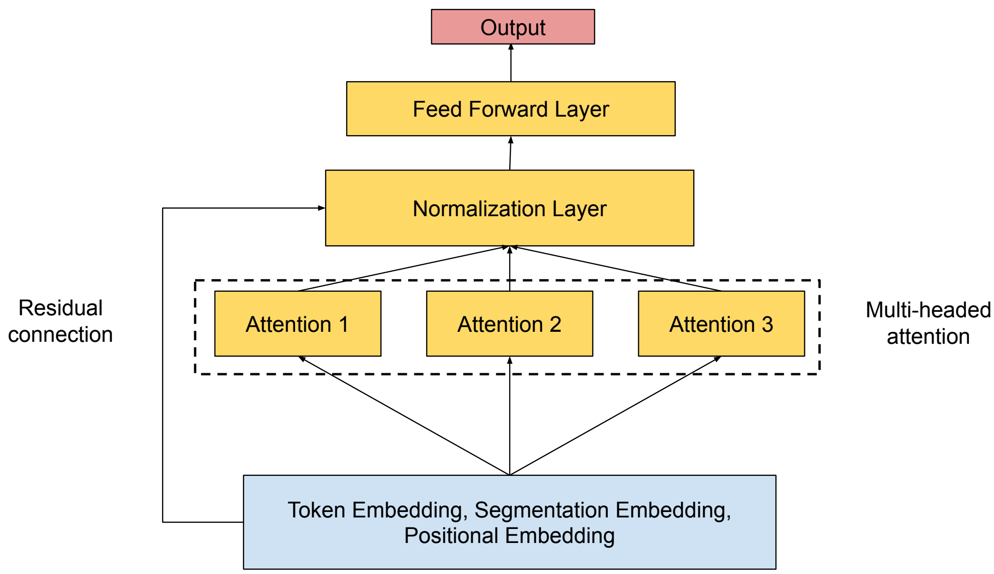

FIGURE IB2: ENCODER MODULE

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.

Appendix C Additional Figures and Tables

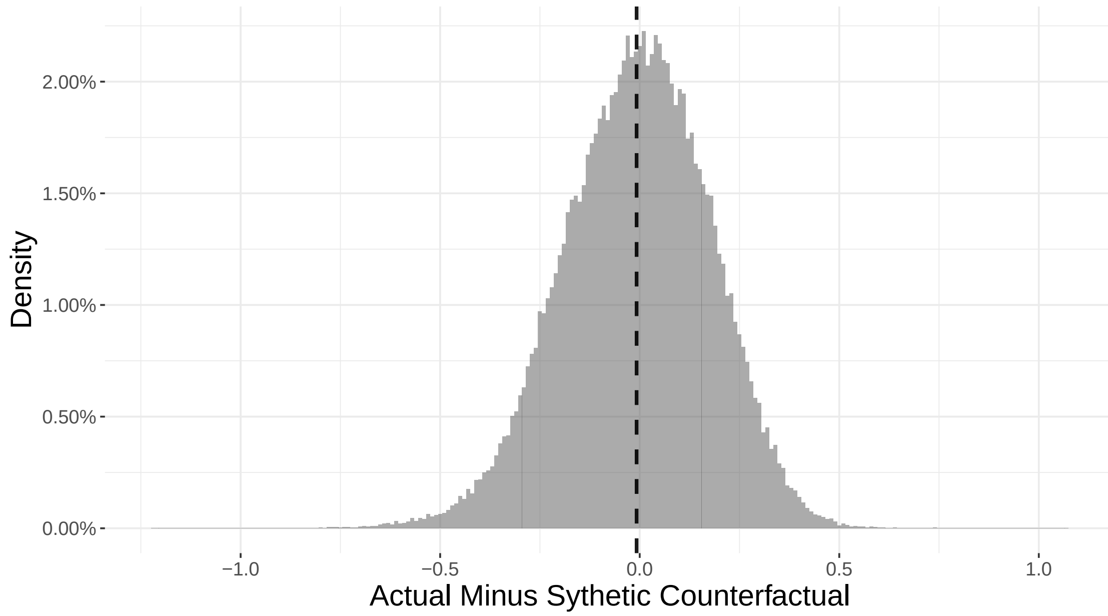

FIGURE IC1: SYNTHETIC DATA TESTS

The figure evaluates the quality of fit of our neural network. The black dashed line is centered at the mean difference in the data (-0.008). We are unable to reject the null that the true difference in means is equal to zero (p=0.2693).

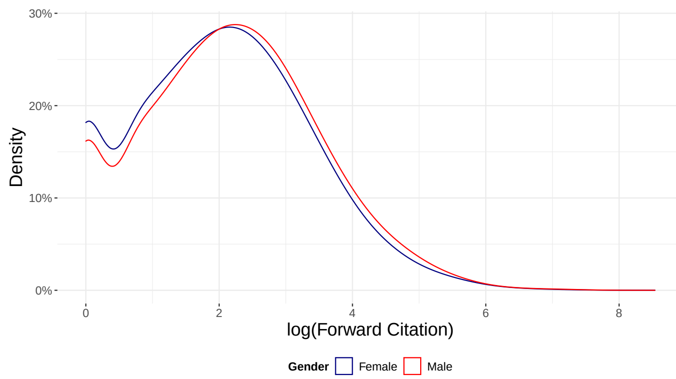

(a) Forward Citations

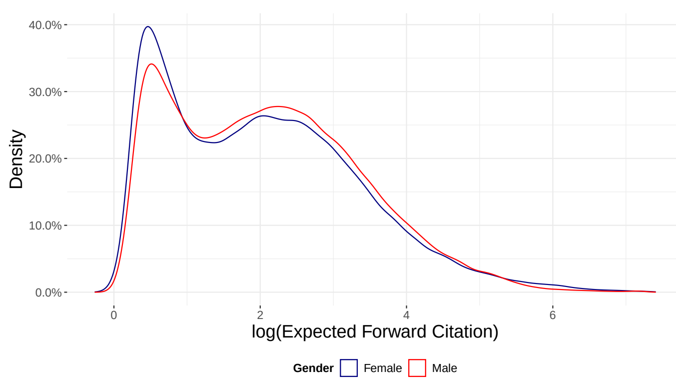

(b) Model Implied Forward

FIGURE IC2: 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.

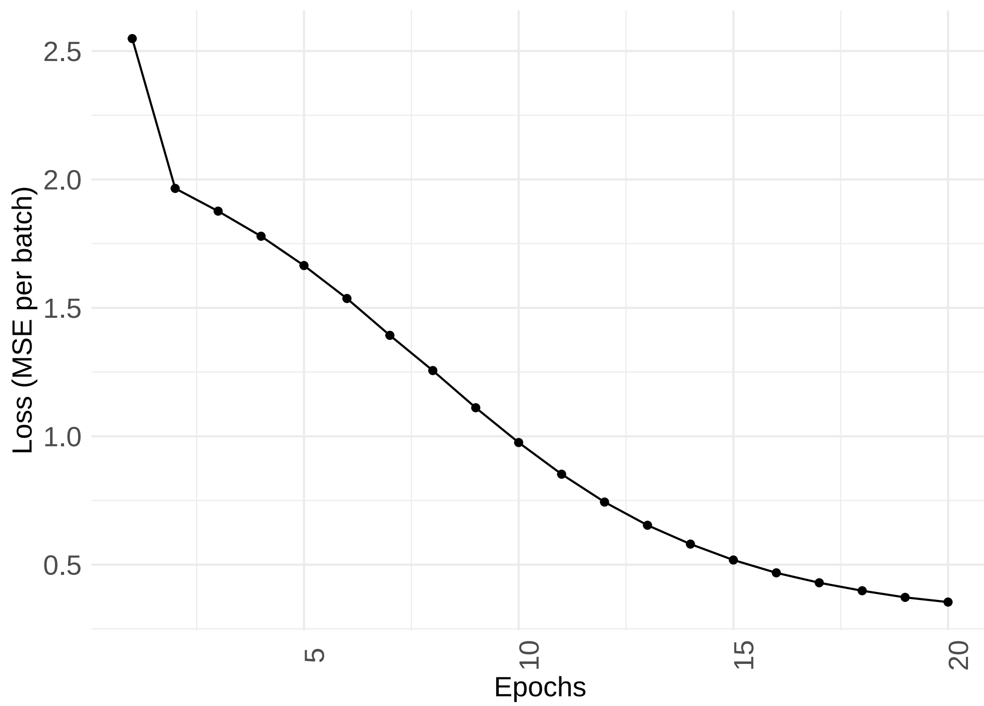

FIGURE IC3: LOSS FUNCTION

This figure illustrates the loss function of the C-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.

TABLE IC1: ROBUSTNESS TO LANGUAGE MODEL

This table replaces the Longformer model with alternative language Models. The sample covers patents issued from 1976-01-01 through 2021-12-31. 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 6. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-3.142***	-3.312***	-3.459***	-3.469***	-3.520***	-3.477***
	(0.158)	(0.157)	(0.148)	(0.143)	(0.135)	(0.139)
Intercept	-3.011***					
	(0.126)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	16.6%	17.4%	18.2%	18.3%	18.5%	18.3%
Observations	774315	774315	774315	774315	774315	774315
R^2	0.002	0.041	0.048	0.093	0.177	0.331

TABLE IC2: CITATIONS IN TOP DECILE

This table studies patents that receive forward citations in the top decile. Panel A documents the relationship between a patent's lead inventor's gender and the propensity to receive citations placing them in the top decile. Panel B documents the relationship between a patent's lead inventor's gender and the model's prediction a patent would be in the top decile of citations. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Forward Citation

	Top Decile Patent					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-0.0169***	-0.0156***	-0.0147***	-0.0140***	-0.0104***	-0.0062***
	(0.0021)	(0.0018)	(0.0019)	(0.0017)	(0.0015)	(0.0015)
Intercept	0.1071***					
	(0.0087)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	17%	15.6%	14.8%	14.1%	10.5%	6.2%
Observations	534997	534997	534997	534997	534997	534997
R^2	0.001	0.088	0.104	0.153	0.245	0.407

Panel B: C-TEXT (Longformer)

	Flipped to Top Decile					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	0.00972***	0.00985***	0.00996***	0.01016***	0.01054***	0.01077***
	(0.00061)	(0.00058)	(0.00056)	(0.00056)	(0.00052)	(0.00057)
Intercept	0.00895***					
	(0.00057)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	9.8%	9.9%	10%	10.2%	10.6%	10.8%
Observations	534997	534997	534997	534997	534997	534997
R^2	0.002	0.006	0.008	0.038	0.114	0.246

TABLE IC3: ROBUSTNESS TO SAMPLE SELECTION

This table establishes the robustness of our baseline specification of Panel B of Table 3. 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 6. Panel A uses a single-author patent. Panel B uses the majority of genders of the inventors. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Single Author

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-2.030***	-2.074***	-2.083***	-2.204***	-2.301***	-2.00***
	(0.275)	(0.258)	(0.259)	(0.259)	(0.248)	(0.26)
Intercept	-0.200					
	(0.134)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	10.5%	10.7%	10.7%	11.4%	11.9%	10.3%
Observations	213540	213540	213540	213540	213540	213540
R^2	0.001	0.022	0.023	0.132	0.272	0.516

Panel B: Majority Same Gender

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Majority Female Inventors	-0.438*	-0.517**	-0.516**	-0.637**	-0.752***	-0.871***
	(0.186)	(0.183)	(0.183)	(0.188)	(0.194)	(0.185)
Intercept	-0.935***					
	(0.177)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	2.1%	2.5%	2.5%	3.1%	3.6%	4.2%
Observations	211931	211931	211931	211931	211931	211931
R^2	0.000	0.017	0.018	0.107	0.262	0.486

TABLE IC4: PLACEBO TEST

This table presents a placebo test by randomizing the gender of patents and re-running our C-TEXT approach to establish the effects are not an artifact of C-TEXT. 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 6. The sample covers patents with a positive number of citations, granted from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-0.0916	-0.0919	-0.0923	-0.1148	-0.0922	-0.0835
	(0.0870)	(0.0878)	(0.0879)	(0.0956)	(0.0892)	(0.1140)
Intercept	0.919***					
	(0.142)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	0.5%	0.5%	0.5%	0.6%	0.5%	0.4%
Observations	238081	238081	238081	238081	238081	238081
R^2	0.000	0.007	0.008	0.089	0.219	0.415

TABLE IC5: ROBUSTNESS TO TRAINING SAMPLE

This table presents robustness tests for the time-period of the sample. 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 6. Panel A is trained and tested on the sample from 1976-01-01 through 1999-12-31, while Panel B uses the sample from 2000-01-01 through 2021-12-31. Standard errors are clustered at the patent attorney and patent issue year level. ***, **, *, + denote significance at the .1%, 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Pre-2000 Sample

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-2.116***	-2.012***	-2.001***	-2.035***	-2.051***	-1.824***
	(0.274)	(0.238)	(0.241)	(0.236)	(0.225)	(0.179)
Intercept	0.931***					
	(0.107)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	7.6%	7.2%	7.1%	7.3%	7.3%	6.5%
Observations	177898	177898	177898	177898	177898	177898
R^2	0.001	0.010	0.010	0.089	0.213	0.460

Panel B: Post-2000 Sample

	Delta					
	(1)	(2)	(3)	(4)	(5)	(6)
Lead Female Inventor	-0.757***	-0.972***	-1.06***	-1.07***	-1.161***	-1.16***
	(0.189)	(0.191)	(0.19)	(0.20)	(0.199)	(0.19)
Intercept	-1.075***					
	(0.261)					
Art Group FE	No	Yes	Yes	Yes	Yes	Yes
Patent Issue Year FE	No	No	Yes	Yes	Yes	Yes
Examiner FE	No	No	No	Yes	Yes	Yes
Attorney FE	No	No	No	No	Yes	Yes
Assignee FE	No	No	No	No	No	Yes
Relative to Sample Mean	4.9%	6.3%	6.9%	6.9%	7.5%	7.5%
Observations	429994	429994	429994	429994	429994	429994
R^2	0.000	0.021	0.023	0.072	0.157	0.315