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

March 2023

Abstract

Implementing a state-of-the-art machine learning technique for causal identification from text data, we document that patents lead-authored by female inventors are under-cited relative to their quality. For the equivalent patent with a lead female inventor, a patent with a male lead inventor would have received 28% more citations. Male lead inventors in particular tend to undercite patents with female lead inventors, while patent examiners of both genders appear to be more even-handed. Market-based measures of patent value load on the citation counts that would be predicted for a female lead inventor patent had it been lead-authored by a male, rather than on actual citation counts. 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, 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, Dong Yang, Tengfei Zhang, participants at the NBER Productivity Innovation and Entrepreneurship Meeting, Midwest Finance Association (MFA), York University (Schulich), and the Mid-Atlantic Research Conference in Finance (MARC) for their helpful comments and suggestions. Yael V. Hochberg is with the Jesse H. Jones Graduate School of Business at Rice University and NBER, email: hochberg@rice.edu. Ali Kakhbod is with the Haas School of Business at the University of California, Berkeley, email: akakhbod@berkeley.edu. Peiyao Li is with the Haas School of Business at the University of California, Berkeley, email: ojhfklsjhl@berkeley.edu. Kunal Sachdeva is with the Jesse H. Jones Graduate School of Business at Rice University, email: 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, 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 measurement of the causal contribution of gender to 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. Here, assuming that the text of a patent contains sufficient information to adjust for confounding (common) causes between the treatment and outcome, we can use textual analysis to identify the causal effect of treatment, as shown in Figure 1.

For our main analysis, we adopt the causal bidirectional encoder representation from transformers or C-BERT (Veitch, Sridhar, and Blei (2020)). C-BERT estimates causal effects from observational text data, adjusting for confounding features of the text, such s the subject or writing quality. It assumes that the text content suffices for causal identification but is prohibitively complex for standard analysis. C-BERT 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 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 information necessary to measure the desired effects (quality of the patent and forward citations, conditional on gender).

Our main sample covers all utility patents granted by the U.S. Patent Office (USPTO) from 1976 through 2021. The extended follow-up period allows us to measure the impact of a patent without concerns of forward citations being right-censored. We textual matching the first names of each inventor to data from the Social Security Administration and the World Intellectual Patent Organization. We then label each patent by the inferred gender of the inventors following Desai (2019). For our main analyses, we 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 labelings of female versus male-authored patents, including restricting to

single-authored patents or gender-homogenous teams.

The next step is to train the model. First, we use a pre-trained BERT model provided by Google (from Python packages Transformers and Tensorflow) to transform the text of each patent's abstract[1] into a numerical representation. Next, given that the sample is dominated by male-authored patents, we reduce the sample to an equal set of male- and female-authored patents by randomly selecting a subsample of male patents without replacement. Finally, we train two neural networks—one per gender of the lead patent author—using the BERT numerical representations as inputs and citations as outputs, representing 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 the 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 patent data for each gender and run it through the prediction 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 lead author.[2]

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). This pattern persists when we control for factors such as the identity of the patent examiner, the identity of the correspondent for the patent submission (usually the law firm involved), the art unit of the

[1] we note that our main results are robust to using full patent text and an appropriate textual embedding approach (LONGFORMER) for longer texts instead of BERT.

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

patent, and the patent issue year. The distributional plots of forward citations confirm the statistical tests and suggest female-authored patents receive 1 to 3 fewer citations than male-authored patents. More generally, patents authored by women appear to be less likely to receive any citations than male-authored patents.

We then use the causal text analysis methodology to mediate for differences in the quality of patents. We perform two sets of analysis. First, we explore the extensive margin, including in our sample all patents, including those that receive zero citations (the modal patent). Second, we explore the intensive margin, restricting our sample to those patents that receive at least one forward citation. In both cases, we control for a variety of patent characteristics and fixed effects.

At the extensive margin, applying the C-BERT methodology and comparing actual citations to those that would be predicted if the patent was authored by an inventor of the opposite gender, we find that patents with a female lead inventor received approximately 28% 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 3.8 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: patents with a female lead inventor received approximately 20% fewer citations than an equivalent patent would receive had the lead inventor been male, a difference again of approximately 3.8 fewer citations per patent.

Overall, the results suggest that patents with a female first author would receive more citations if their first author had been male. The results are robust to a variety of alternative specifications and do not seem to be attributable to sample selection or model overfitting. The results are also robust to a variety of different 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 only of female authors versus patents authored by teams composed only of male authors. Finally, while our main analysis uses patent abstracts, this result is also robust to employing the full text of the patent instead of the abstract.[3]

Our results hold across NBER's major categories and subcategories of patent technology, with some heterogeneity in subcategories.We observe similar patterns when using the Cooperative Patent Classification (CPC). We observe similar differences in citations in both emerging and established technology fields. This undercitation of female-authored patents grows over time, becoming particularly more pronounced since the early 1990s.

We then turn to explore the source of these patterns. Under-citation of female-authored patents could be driven by the citations inventors themselves include in their patent applications, or by citations that are added by patent examiners. Controlling for art unit, issue year, and examiner, patents with female and male first authors both significantly undercite patents with a female first authors. Patents with male first authors, however, do so by approximately double the margin. In contrast, we see little evidence of either male and female examiners underciting female-led patents. Overall, the results suggest that the undercitation of female patents is largely due to patent with male first authors underciting past female-authored patents in their patent applications.

Finally, we explore whether the expected value accorded to the patent by the market reaction to its issuance correlates more with actual citations or with the citations predicted by 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. We regress these measures of economic value on our adjusted forward citations, actual forward citations, and a variety of controls. When horseraced against each other, measures of expected economic value load significantly on our adjusted forward citation measure,

[3]Due to the differences in text length between patent full text and patent abstract text, using BERT for the embedding is computationally infeasible. Instead, we utilize an embedding called LONGFORMER, which is more suitable for longer texts.

but not on actual forward citations. The results suggest that expected economic value, as measured by market reactions, may be 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 patent text used captures all the factors that should influence the number of forward citations. While it is not possible to test this assumption directly, it is reasonable to assume that the text of the patent is closely related to its quality or importance. Second, the predictive ability of any given model in computing counterfactual outcomes is difficult to assess in these settings. Instead, we plot the loss function of our model training and find that the fitting error starts to plateau after around 15 epochs of training. Therefore, in order to prevent overfitting, we stop training shortly afterat 20 epochs. Note that the remaining error is an order of magnitude smaller than the initial error, so the model has made considerable progress in fitting the data. Further, we examine the quality of our model fit by synthetically creating counterfactual data to evaluate the accuracy of our model. As shown in Figure IA4, we find a high accuracy of the model without bias. 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).[4] 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. Moreover, this may exacerbate the already substantial wedge between men and women in science, technology, engineering and mathematics (STEM) fields (Beede, Julian, Langdon, McKittrick, Khan,

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

and Doms, 2011), leading to further inefficient allocation of labor. A second implication concerns 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 in economics, finance, and innovation that relies on forward citations to measure patent quality, these findings suggest that a re-examination of prior findings may be warranted.

Our findings contribute to an emerging literature studying obstacles that inventors face in the U.S. patent system. Research has studied 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 all 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 causally identifying differences in patent applications and approvals, our findings focus on a relatively unexplored question: whether women also face obstacles in citation of their patents.

Relatedly, our findings 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 Card, DellaVigna, Funk, and Iriberri (2020) and Hengel and Moon (2020) show that, controlling for quality, female academics in economics receive fewer citations for their work. Koffi (2021) 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. Chawla (2016) and Koffi and Marx (2021) study broader academic fields. Our work suggests parallels

in patent citations as well.

In addition, our paper makes an important methodological contribution. A large literature across many fields has demonstrated that big data has the potential to revolutionize research in general and finance and economics research in particular (Goldstein, Spatt, and Ye (2021)). In economics, a small but rapidly growing branch of the big data literature uses natural language processing to quantify text, allowing it to be used in empirical applied microeconomics research (Gentzkow, Kelly, and Taddy (2019a)). 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). Recent advances in computer science have produced new methods that allow the use of text embedding to mediate and identify causal effects. Our paper introduces these methods to finance and economics, proposing a new technique, C-BERT, that uses text as a mediator, allowing us to causally identify the effect of gender on citations of patents. To the best of our knowledge, we are among the first researchers to apply deep learning in economics and finance 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 granted by the U.S. Patent Office (USPTO) from 1976 through 2021. This allows for at least a 20-year follow-up history, extending through the patent's expiration.

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 abstracts, as they provide a clear and concise text-only summary of the core contribution of the patent. Importantly, this is the key text input into the C-BERT model. Using the abstract of patents presents several key advantages over using the full body. First, use of abstracts alleviates concerns about differences in the quality of the figures contained within the patents that could substitute for the quality of the writing. Second, abstracts are a good proxy for the contents of a patent as well as what inventors and examiners review. Third, from a practical standpoint, using the full text of the patents is computationally prohibitive. Even with access to a high-powered computing cluster, using abstracts in our setting takes several days to complete.

In robustness tests, however, we reproduce our analysis using the full text of the patents, and substituting the LONGFORMER embedding for BERT, as BERT does not scale for long texts. All our results remain qualitatively similar (we intend to include these results in more detail in a future draft).

1.2 Patent Citations

Importantly, there is a key difference between patent citation counts and actual patent quality. While historically forward citations have been used as a proxy for patent quality, the key point of our analysis is to determine whether this measure is systematically biased downwards for female-authored patents. We therefore distinguish between patent citations (the easily observed outcome for a patent) and quality, which is mediated for by using the text of the patent. Patent forward citation counts are obtained through use of data from the USPTO.

Of course, 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.

1.3 Gender of Inventors and Examiners

Our main treatment variable is the gender of the lead inventor (first author).[5] Inventors, however, do not disclose their gender when applying for a patent. 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 the first name of the lead inventor to identify the gender of the inventor (Tzioumis, 2018).

Starting with the PatentView data, we obtain the first names of each inventor of each patent. For patents with multiple inventors, we rely on the name of the first inventor due to that person's prominence. Next we classify the gender of patent inventors 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 utilizing a cross-country dataset from the World Intellectual Property Organization (WIPO) (Martinez, Raffo, Saito, et al., 2016). We drop patents when there is no distinct gender determination for the lead inventor.

One challenge is that our sample shows that women are underrepresented as inventors on patents (Hunt et al., 2013). As a result, we need to balance our sample across patents with lead inventors from each gender. To do this, we use all patents with a female lead inventor and extract a random subsample of patents with male lead inventors of the same size. We then estimate a propensity model using a one layer logit-linear

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

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

neural network, where the objective function is the binary-cross-entropy between the predicted treatment indicator and the true treatment indicator. The output of this neural network is the predicted probability that this patent is written by a female lead author. Following the C-BERT methodology in Veitch, Sridhar, and Blei (2020), we then drop (i) all patents in the male subsample whose estimated propensity of being female-authored is very low (less than 3%) and (ii) the all patents in the female subsample whose estimated propensity of being female-written is very high (greater than 97%). This allows for any given patent in the remaining sample to have some level of ambiguity, based on textual analysis, as to whether it was authored by a male of female.[7]

1.4 Examiner and 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 were explicitly added by examiners and inventors.

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

[7] All our results remain qualitatively similar in nature and stronger in magnitude if we do not exclude patent whose author gender can be clearly identified from the text content alone.

proxy for technology type grouping. Our baseline sample contains 898 art units and 11,953 patent examiners.

As an alternative and intuitive patent technology grouping is provided by the NBER patent category, which is also reported in the USPTO PatentView database.[8] The NBER classification includes six major categories (computer and communications, drugs and medicine, electrical, mechanical, chemical, and other) and 37 sub-categories. We use the six categories and 37 sub-categories to examine heterogeneity by patent technology type, allowing us to present individual subcategory estimates in a digestible manner. In further robustness tests we also consider the Cooperative Patent Classification (CPC) of the patents, (see Table IA3).

Patents are typically filed with the assistance of a patent attorney, who may file many them on behalf of different inventors. The USPTO refers to these law firms (or legal department of the patent assignee firm) as "customers," and each such entity is assigned a customer number. Approximately 60% of observations have a legitimate customer number. If the observation lacks a valid customer number, we assign it a common value ("unassigned"). These identifiers are useful because they allow us to account for possible commonalities in writing style across patent attorneys that may influence the text of the final submission. Our baseline sample contains 9,516 unique customers. We label these customer identifiers as "law firm."

Descriptive statistics for our sample are presented in Table 1.

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

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

citations. Finally, we must compute the counterfactual of citations based on the gender of the inventor.

Below, we outline our empirical strategy to address these challenges. 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-BERT 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. Here we use a Bidirectional Encoder Representations from Transformers (BERT) approach to transform the text of each patent 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 constructs embedding vectors that are numerical representations of the text, which preserve both the meaning of individual words and the underlying context of each word.[9] The BERT encoder module (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, and Polosukhin (2017)) produces a high-dimensional representation with 768-dimensional embeddings that each represent the text of a patent's abstract. We describe the encoder architecture in detail in Appendix A.

2.2 Causal BERT (C-BERT)

Having created high-dimensional representations of patent text, the second challenge is to establish the relationship between this data and patent citations. To do so, we use a novel leading machine learning technique called Causal Bidirectional Encoder Representations from Transformers (*C-BERT*) that allows us to causally estimate the contribution

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

of language on a binary treatment variable. C-BERT comes from recent advances in computer science, including Khetan, Ramnani, Anand, Sengupta, and Fano (2022); Shao, Li, Gu, Qian, and Zhou (2021), which has developed methods to use text embedding as a mediator Veitch et al. (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. To the best of our knowledge, ours is among the first papers to apply deep learning to causal inference with language in the field of economics and finance.

C-BERT is a neural-network-based architecture that estimates counterfactuals of a binary treatment under the assumption that all of the information (covariate) needed for causal identification is contained within a given text. As shown in Figure IA2, the input data for training contains three types of information: the texts of 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 BERT model for generating text embeddings, a logit-linear model the 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.

The C-BERT methodology in our context has two key steps. First, it uses a pre-trained BERT embedding to transform the text of each patent into a high-dimensional numerical vector. The embedding vectors are numerical representations of the text that preserve both the meaning of individual words and the underlying context of each word. Second, C-BERT computes 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? We estimate this by passing the subset of patents with a male lead inventor through the mapping trained on the embedding vector of female lead inventor patents, and passing the subset of patents with female lead inventors through the mapping trained on the embedding vector of male-lead-inventor patents. From this, our mapping estimates the number of citations for the counterfactual number of citations by gender.

The procedure is depicted in Figure 2. First, we run the trained C-BERT 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 BERT 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.

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 BERT model f to map W_i to Z_i where Z_i is the embedding of the abstract. 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.

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

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:

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

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.

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

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

The resulting output of our C-BERT model is the actual outcome and a counterfactual outcome. In our application, this is the number of citations and the estimated number of citations the opposite gender would have received.

2.3 Assessing C-BERT's Identification Assumptions

There are three assumptions that the econometrician must consider when applying C-BERT.

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 quality of the patent should be measurable by the content (text) of the patent itself. Patent examiners read the text of the proposals to evaluate the novelty of patents prior to granting a patent. As a result, this necessary condition is likely 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.

To assess the quality of our embedding representations, we consider synthetic tests to measure the accuracy of our model. To do this, we first compute the synthetic outcomes of all of the patents across the full dataset. In doing this, we used a random linear transformation that takes a uniformly random 768 vector with values from 0 to 1. Then we take the dot product of this random vector with each patent's 768 dimensional embedding. Finally, the resulting values are the synthetic outcome for females, and, for males, we add a known scalar to the function. In this approach, we know the true treatment

effect and can evaluate the model. Applying the C-BERT model uncovers the known true treatment effect with an accuracy of over 90%. This high level of accuracy suggests that the embedding method clearly extracts semantically meaningful information.

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 forward citation counts by gender using simple regression analysis. We then apply the C-BERT model to calculate counterfactuals, and assess the causal effect of gender on citation counts.

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 female receive fewer citations than males.[10] On average, male lead inventors receive significantly more citations than female lead inventors (18 citations for males, 15 for females, F-stat = 518, Table 1). Testing the difference in distributions, we find a Kolmogorov-Smirnov statistic of $D = 0.066426$, 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 estimating the following OLS model:

$$Y_i = \beta_1 I\left(FemaleInventor_i\right) + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Customer \times Examiner} + \varepsilon_i, \tag{1}$$

where patent and year are represented by i and t, respectively. Y_i is our outcome of interest. Our specification includes fixed effects for customer-examiner pair ($\delta_{customer \times examiner}$), 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 customer. β_1 is our coefficient of interest, where a positive value would indicate that women receive more citations than males, and vice versa.

The estimates are presented in Table 2. Panel A presents the 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).[11] Both panels suggest that female lead investors receive between 0.8 to 4 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

[10]To address skewness and show more clearly, we the natural logarithm of citations and present a histogram in Panel A of Figure IA1.

[11]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.

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 Adjusting for the Quality of Patents Using C-BERT

To explore this, we turn to our C-BERT model. As a reminder, C-BERT 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 citations a patent would have received had its lead author been of the opposite gender, $\widehat{ForwardCitation_i}$. We plot the histogram of predicted citations from C-BERT by gender in Panel B of Figure 3.[12]

We compare the actual forward citations to the model-implied citations at the patent level. Specifically, we calculate the difference in actual versus predicted citations as:

$$Delta_i = ForwardCitation_i - \widehat{ForwardCitation_i}, \tag{2}$$

where $ForwardCitation_i$ is the actual number of citations to a given patent authored by a given gender and $\widehat{ForwardCitation_i}$ is the number of citations implied if the lead inventor had been of the opposite gender. A positive delta implies that a patent has received more citations than the quality-adjusted number suggested by the opposite-gender model.

We plot the difference between actual and model implied citations in Figure 4 for patents with at least one citation. We observe that $Delta_i$ appears to be negative on average for female lead inventors (plotted in red), with a mean and median of -2.69 and

[12]To address skewness and to show the distribution more clearly, we the natural logarithm of citations and present a histogram in Panel B of Figure IA1.

-0.25, respectively. In contrast, for male lead inventor patents (plotted in blue), we see a positive delta versus what they would have been predicted to receive with a female lead author instead, with a mean and median of 1.10 and 0.13, respectively. For both genders, these difference measures show evidence of skewness. Notably, it is clear that the delta for any given patent is not strictly positive for males or negative for females.

To study the relative delta, we re-estimate the OLS in Equation 1, replacing the actual number of forward citations for a patent with $Delta_i$. The estimates are presented in Table 3. Panel A presents estimates for the full sample of patents, including those with zero citations (extensive margin), while Panel B presents the estimates for the subsample of patents which 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 otherwise the same, except that the lead inventor was instead male. Interpreting our point estimate relative to the sample mean of the number of forward citations, we find that patents with female lead inventors receive 3.8 fewer citations than would be predicted if they had had a male first author instead, controlling for other characteristics and mediating for quality using C-BERT. The magnitudes are relatively unchanged when including patent-year fixed effects, art-unit fixed effects, and various customer and examiner fixed effects. The relative stability in estimates suggests that our analysis does not suffer from a correlated omitted variable, Oster (2019).

Next, we consider the difference in the propensity for female lead patents to be in the top-decile of patents by forward citations. To study this we once again estimate Equation 1 but replace the dependent variable with an indicator that takes the value of one if the patent is in the top decile of citations based on the mapping for the opposite gender but is not based on observed data. The estimates are presented In Panel A of Table 4. The estimates suggest that without any model adjustments, female lead inventor patents are underrepresented in the top decile of patents by citations, relative to patents

with a male first author.

We then apply the C-BERT methodology adjustment. In Panel B of Table 4, we re-estimate the models, where the dependent variable takes the value of one if the patent based on actual citations was not in the top decile of citations, but under the C-BERT predicted forward citation, would be ("flipped to top decile"). The estimates suggest that the probability of a patent flipping into the top decile under the counterfactual is approximately 1.7 percentage points higher if the lead inventor is female. These estimates are highly statistically significant and economically meaningful.

These three sets of results represent causal evidence suggesting that female lead inventors are undercited, on average, relative to an equivalent patent granted to their male counterpart. These differences in forward citations cannot be explained by differences in art units, time trends, or differences in customers or examiners.

4 Cross Sectional Heterogeneity

Next, we explore whether these patterns of undercitation are uniform across a variety of dimensions of heterogeneity in patent characteristics.

4.1 Patent Category

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.

$$\begin{aligned} Y_i &= \beta_1 I\left(FemaleInventor_i\right) + \beta_2 I\left(FemaleInventor_i\right) \times \left(Category\right) \\ &+ \delta_{PatentCategory} + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Customer \times Examiner} + \varepsilon_i, \end{aligned} \tag{3}$$

where the subscript and notation match the prior estimating equations. As in the main analysis, standard errors are double clustered by year and customer. First, we interact our female indicators with the six NBER categories to study differences by broad field categories. Then we explore the 37 subcategories in a similar manner.

The estimates in Table 5 highlight important heterogeneity across patent categories. Column (1) of Table 5 presents the estimates for the model using the major categories, 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 implied by the C-BERT model for the opposite gender, *Delta*. To interpret the overall effects for our variable of interest for each category, we need to add the coefficients of the indicator for female lead inventor with the interaction term for each major category. In general, we observe some level of undercitation for all patent categories, with particularly large disparities for the Drugs and Medical category. Put differently, if a female lead-inventor patent instead had a male lead inventor, it would have received significantly more citations, regardless of technology category, echoing our baseline results.

We can further break down the technology category using the 37 NBER subcategories. Specifically, we estimate the following.

$$\begin{aligned} Y_i &= \beta_1 I\left(FemaleInventor_i\right) + \beta_2 I\left(FemaleInventor_i\right) \times \left(Subcategory\right) \\ &+ \delta_{PatentSubcategory} + \delta_{GrantYear} + \delta_{Customer \times Examiner} + \varepsilon_i, \end{aligned} \tag{4}$$

To ease interpretation, Figure 5 presents the sum of the female lead indicator and the interaction coefficients (Female Lead Inventor $\times$ Subcategory) graphically. The raw estimates are presented in the Internet Appendix in Table IA4.Column 1 of Table IA4 shows the estimates employing the raw citation counts as the dependent variable, while column (2) uses *Delta* as a dependent variable. 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 in the main analysis, standard errors are double clustered by year and customer. As can be seen clearly in Figure 5, for the vast majority 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. 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 five 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 6. While our main result is still apparent, with patents with female first authors exhibiting an estimated 3.3 to 3.7 fewer citations than would be predicted if the first author had been male, depending on specification), across all specifications, we cannot reject the null that there is no additional citation difference

in emerging versus established fields.

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 diminish 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 subperiods, for each patent, we collect the forward citations the patent receives during this sub period post grant. For each subperiods, we re-run our C-BERT methodology to estimate the delta in citations after mediating for patent quality. This allowed us to examine when undercitation occurs, relative to the time of patent grant.

The estimates are presented in Table 7. Column (1) presents estimates from forward citations to patents received in the first year after patent grant, column (2) presents estimates for forwward 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-BERT using only forward citations received during that subperiod (by necessity, the number of observations is smaller in later subperiods as fewer of the patents in our sample will yet have histories of that length). 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 patent grant, consistent with undercitation being self-reinforcing over time. The coefficient on lead female inventor 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 -0.7, -1.3, and -2.7 for the [1-5) years, [5-10) years, and [10-20] years periods, respectively.

Our estimates in Table 7 are also consistent with our baseline specification. When we sum the point estimates across the four periods, we obtain results that are similar in magnitude to those in the baseline specification. Furthermore, when we consider the magnitude of the estimates relative to the sample means, we see that the bias in citations increases monotonically over time, with values of 3%, -17%, -20%, and -24.25% for the [0-1) years, [1-5) years, [5-10) years, and [10-20] years periods, respectively. These findings suggest that biases in citations may be reinforced by prior biases, leading to a situation in which overcited patents continue to be overcited and the bias becomes larger over time.

4.4 Evolution of Biases 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 within the sample, as gender norms and female participation in the workforce more generally and in science and engineering more specifically have been changing over time.

Of course, older patents tend to naturally receive more citations. Without adjustments to our initial methodology, our findings may incorrectly suggest a decrease in bias over time, when in reality, it is simply a reflection that newer patents receive fewer citations on average. To accurately study the evolution of bias, we restrict our measure

of forward citations to the first ten years post patent grant. This method avoids the right censoring problem of forward citations and provides a clearer interpretation, at the cost of excluding forward citations made after ten years out, where we know more of the undercitation activity occurs. Because we need to be able to measure ten years of forward citations, we by necessity must exclude patents granted in the last decade of our sample. We thus create a sample of patents from 1976 to 2011, and measure the number of forward citations they receive within 10 years. Then, we apply our C-BERT technique to calculate the bias for each patent.

We estimate models of the following nature:

$$\begin{aligned} Y_i &= \beta_1 I\left(FemaleInventor_i\right) \\ &\quad + \sum_{j=1977}^{2011} \beta_j I\left(FemaleInventor_i\right) \times I\left(GrantYear = j\right) \\ &\quad + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Customer \times Examiner} + \varepsilon_i, \end{aligned} \tag{5}$$

where β_1 estimates the average undercitation of females across the entire sample, and the set of coefficients β_j estimate the marginal bias in each patent grant year, with 1976 as the year of comparison. Standard errors are clustered in this specification by the grant year. The omitted group is 1976.

Figure 6 plots the interaction coefficient of for each year from equation Equation 7. The time-invariant estimate for coefficient on the female lead inventor variable is -0.34; the interaction coefficients presented in the figure are additive to that number. From the figure, we observe clearly that the average undercitation of patents with female lead authors has become more pronounced over time. In comparison to patents from 1976, those from the late 1970s and early 1980s seem to have been only modestly additionally undercited. However, starting in the 1990s, the additional undercitation of female-led patents rises to around 2 citations per patent. Thus, despite a decrease in disparities and representation of women in the workplace and in science and engineering professions,

undercitation of patents with female first authors seems to be growing over time.

5 Who Undercites Female Inventors?

So far, we have presented causal evidence that patents with female lead inventors receive fewer citations than the equivalent patents with male lead inventors. Next we explore the source of the under-citation: 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 also 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 under-citation, 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 under-citation 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-BERT model, estimating a neural net for male-lead inventor patents and a neural net for female-lead inventor patents to predict forward citation counts by examiners of each gender based on the gender of the lead inventor on the patent of interest. We then run female lead inventor patents through the male neural net model, and vice-versa, to calculate the C-BERT adjustment to mediate for the quality of the patent.

Table 8 presents the results of estimation of regression models using the C-BERT adjustment 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 contribute minimally to the under-citation of female lead inventor patents. The coefficient estimates range from -0.014 citations to -0.029 citations, with only two specification statistically significant at conventional levels. Panel B similarly shows no real evidence of undercitation of female lead inventor patents by male examiners. The coefficient estimates range from -0.006 citations to 0.009, with no specification being statistically significant. Taken together, the estimates suggest that the undercitation bias we observe for female lead inventor patents is not driven primarily by examiner patents.

Having established this fact, we then turn to forward citations added by future inventors to their patent applications. We conduct a similar analysis to that which we conduct above with examiners. The estimates from the C-BERT adjustment regressions are presented in Table 9. 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 suggest a clear pattern. First, as can be seen in Panel A, female lead inventors contribute only modestly to the undercitation of female lead inventor patents. The estimates from the regression models in columns (1) through (4) suggest that female inventors undercite female lead inventor patents by approximately 0.5 citations. In

contrast, the contribution of male inventors to the undercitation of female lead inventor patents is more considerable. The estimates in Panel B suggest undercitation of female lead inventor patents by male inventors by more than 1.3 citations. This is large, both economically and statistically, especially in comparison with our main effect.

Taken together, these results suggest that the undercitation of female lead patents is primarily driven by male lead inventors. Note, however, that we cannot conclude that this is necessarily discrimination on the part of male examiners and inventors. For example, these results may stem from 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. These familiarity networks could be boosted by the presence of a female lead inventor on the current patent. Future research may be necessary to fully distinguish the reason for the underciting.

6 Robustness Tests

6.1 Definition of "Female Authored Patent"

One potential explanation for our findings is that the way in which we classify patents by gender may have spurious influence the results. In our baseline method, we used 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 gender.

To address the possibility that our definition of author gender spuriously produces the patterns we observe, we show robustness of the estimates to a number of alternative approaches to attributing author gender to a patent. First, we limit our sample to patents with only one author, and re-running our C-BERT model, 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.[13]

Using the single-author sample, we still find a statistically significant bias against female inventors. The estimates shown in Panel A of Table IA5 are consistent with those in Table 3. As an alternative, we construct a new sample that includes only patents with all inventors of the same gender, both single authors and teams.[14] When we re-run our C-BERT model using this expanded sample, we find similar patterns of under-citation, with similar coefficient magnitudes (4 versus 3.8). While do not present estimates from our additional analyses using these two samples, they are qualitatively similar to those obtained using a first author gender definition.

6.2 Full Patent Text and LONGFORMER embedding

The results presented up to this point utilize patent abstracts and the BERT embedding. A natural concern is that the patent abstracts do not have enough content to fully pick up patent quality for mediation purposes. For robustness, we repeat our analyses replacing the patent abstract texts with the full patent texts, and utilizing the LONGFORMER embedding instead of BERT. We use the LONGFORMER embedding instead because BERT has difficulties handling longer text lengths.

Table IA6 reports the main results using full patent texts, for the earlier subperiod of our sample, 1986-1994. We observe qualitatively similar, statistically significant, under-citation results (2 citations). In future versions of the paper, we intend to extend this analysis to the full sample period and for all of our estimations.

[13]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 is quality as indicated by text content as opposed to writing style.

[14]Note, this only grows our sample by roughly 2% because teams of all female inventors that can be confidently identified in the sample are rare.

6.3 Overfitting of Model

A standard concern with these types of models is overfitting to the training data. In our setting, we train two different models by completing multiple passes of our training dataset through our algorithm, an epoch. While numerous passes of the data help improve the predictive probability of the neural networks, we could have overfit our model to the data. If so, this would result in relatively poor out-of-sample performance. In the context of our paper, this would result in incorrect or biased out-of-sample predictions of the number of citations.

We address this concern by studying the loss function, as presented in Figure IA5, to ensure a reasonable number of training iterations. 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. Second, we find diminishing improvements to the error rate as we approach 20 epochs. Taken together, these findings suggest that our model is unlikely to be overfitted and, as a result, that the model is appropriate and that reasonable counterfactual citations are predicted 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 and become more pronounced as the years go on. In contrast, the 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 whether these forward looking market estimates of a patent's economic value relate more closely to actual forward citations, or to the predicted number of forward citations we obtain out of C-BERT, which adjusts for author

gender.

To investigate the relationship between economic value and citations, we used 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 fi the first author had been of the opposite gender, loads on one of these measures in particular versus the other.

Specifically, we estimate:

$$\begin{aligned} Y_i &= \beta_1 I\left(FemaleInventor_i\right) + \beta_2 ForwardCitation_i + \beta_3 \widehat{ForwardCitation}_i \\ &\quad + \delta_{GrantYear} + \delta_{ArtUnit} + \delta_{Customer \times Examiner} + \varepsilon_i. \end{aligned} \tag{6}$$

The estimates are presented in Table 10. Only the coefficient on $\widehat{ForwardCitation}_i$ loads significantly, suggesting that the market does not appear to undervalue female-authored patents relative to what it would had that same patent been authored by a male lead inventor. The estimates provides further support to the notion that actual measures of forward citations are biased by gender of the inventor.

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-led patents, relative to equivalent male-led patents.

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 inventors are under-cited relative to male peers with equivalent patents, and their compensation for their 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), leading to inefficient allocations of labor.

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-BERT 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, Goldstein, Spatt, and Ye (2021). 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.

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.

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.

Goldstein, Itay, Chester S Spatt, and Mao Ye, 2021, Big data in finance, *Review of Financial Studies* 34, 3213–3225.

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, 2021, Innovative ideas and gender inequality, Working paper.

Koffi, Marlène, and Matt Marx, 2021, Cassatts in the attic, Working paper.

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.

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.

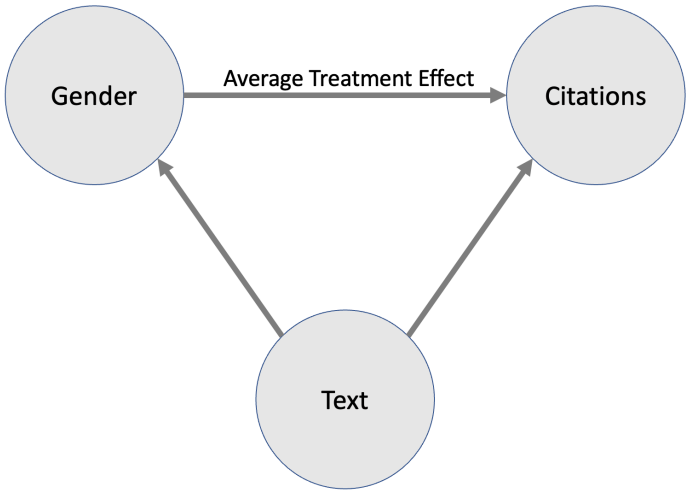
Shao, Yifan, Haoru Li, Jinghang Gu, Longhua Qian, and Guodong Zhou, 2021, Extraction of causal relations based on sbel and bert model, *Database* 2021.

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

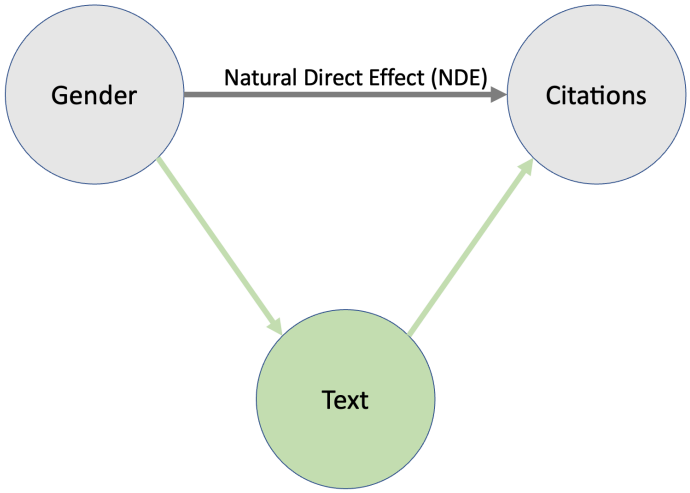
Tzioumis, Konstantinos, 2018, Demographic aspects of first names, *Scientific Data* 5, 1–9.

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: DISTRIBUTION OF FORWARD 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.

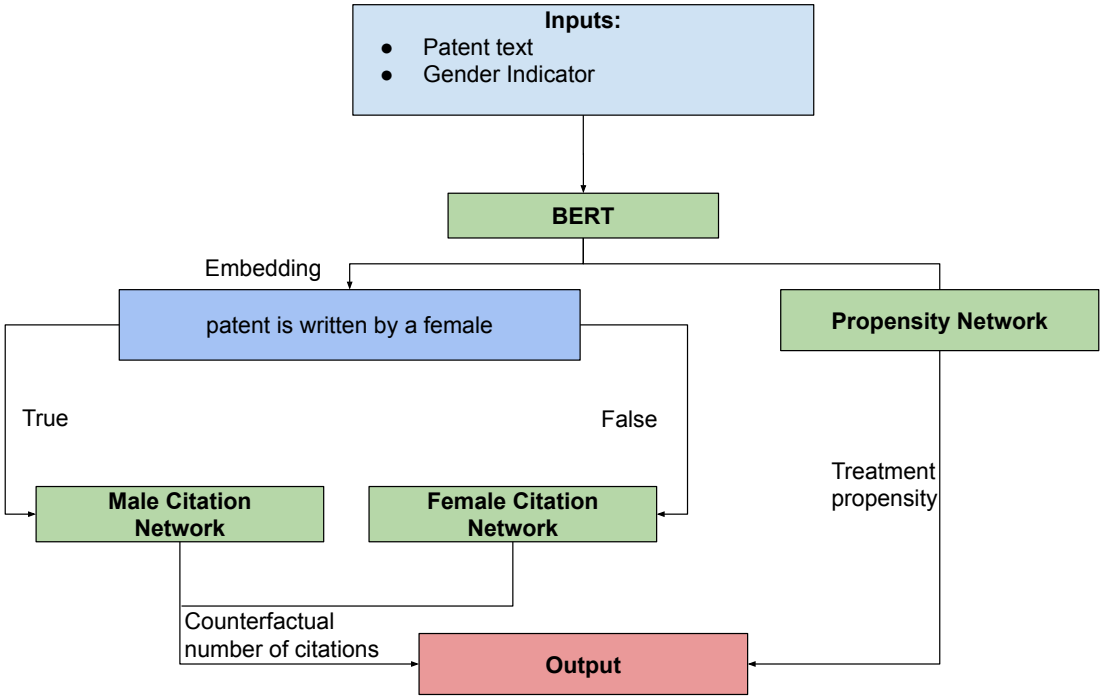
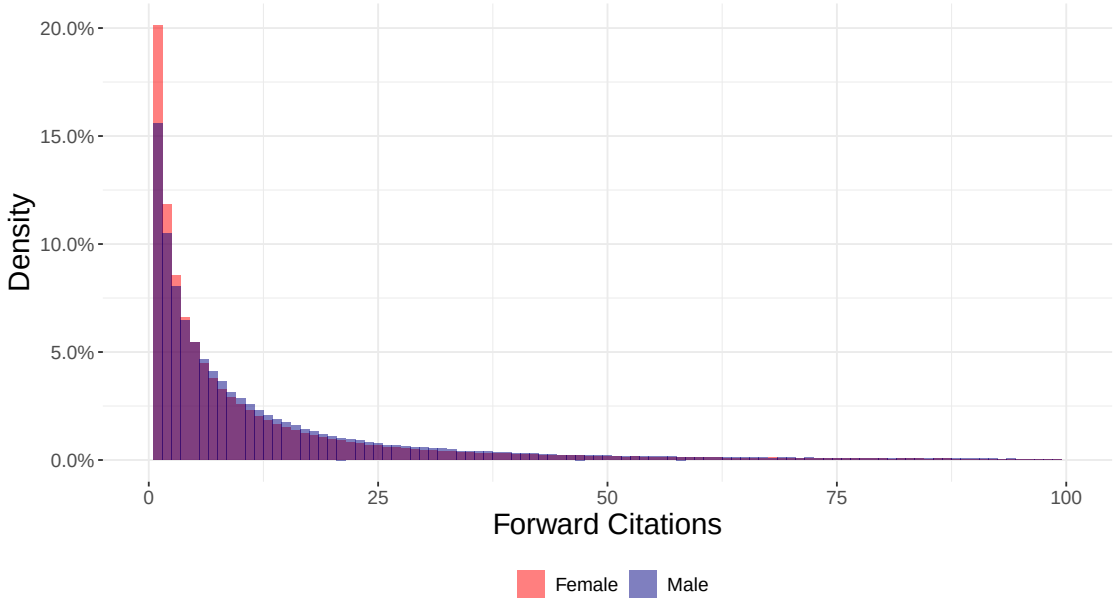
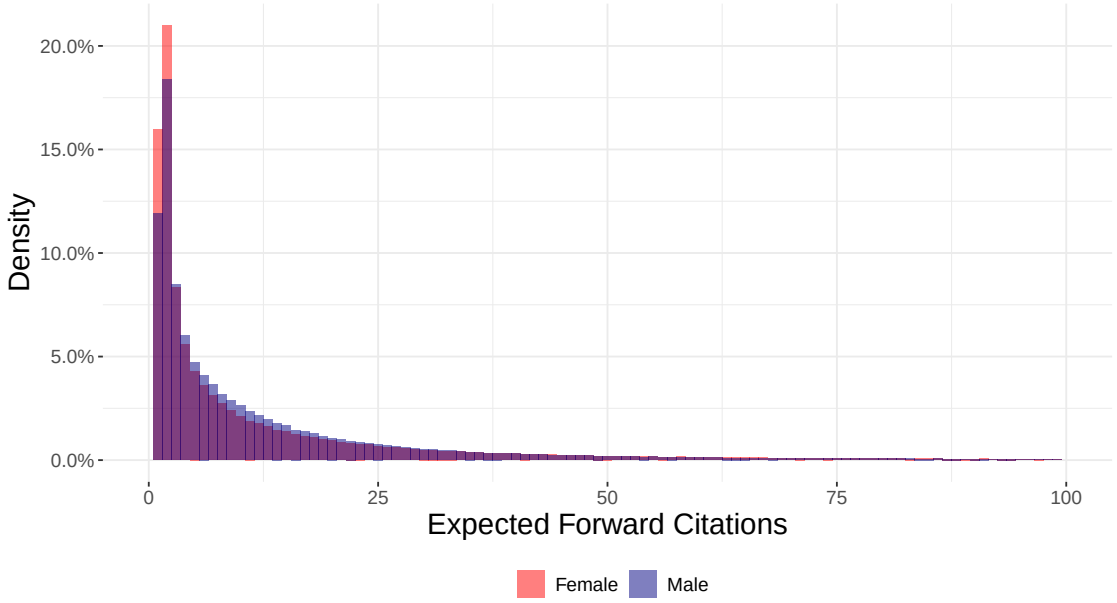


FIGURE 2: C-BERT ESTIMATION PROCEDURE

The figure illustrates the estimation procedure of C-BERT once the neural networks are trained. The light blue block at the very 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: Model Implied Forward

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 the model. The horizontal axis counts the number of citations while the vertical axis measures the percent of the distribution. Red bars correspond to females, blue bars correspond to males, and purple bars correspond to the overlapping region. The distribution is truncated at 100 for ease of interpretation. The natural logarithm transformation of these distributions is presented in Figure IA1.

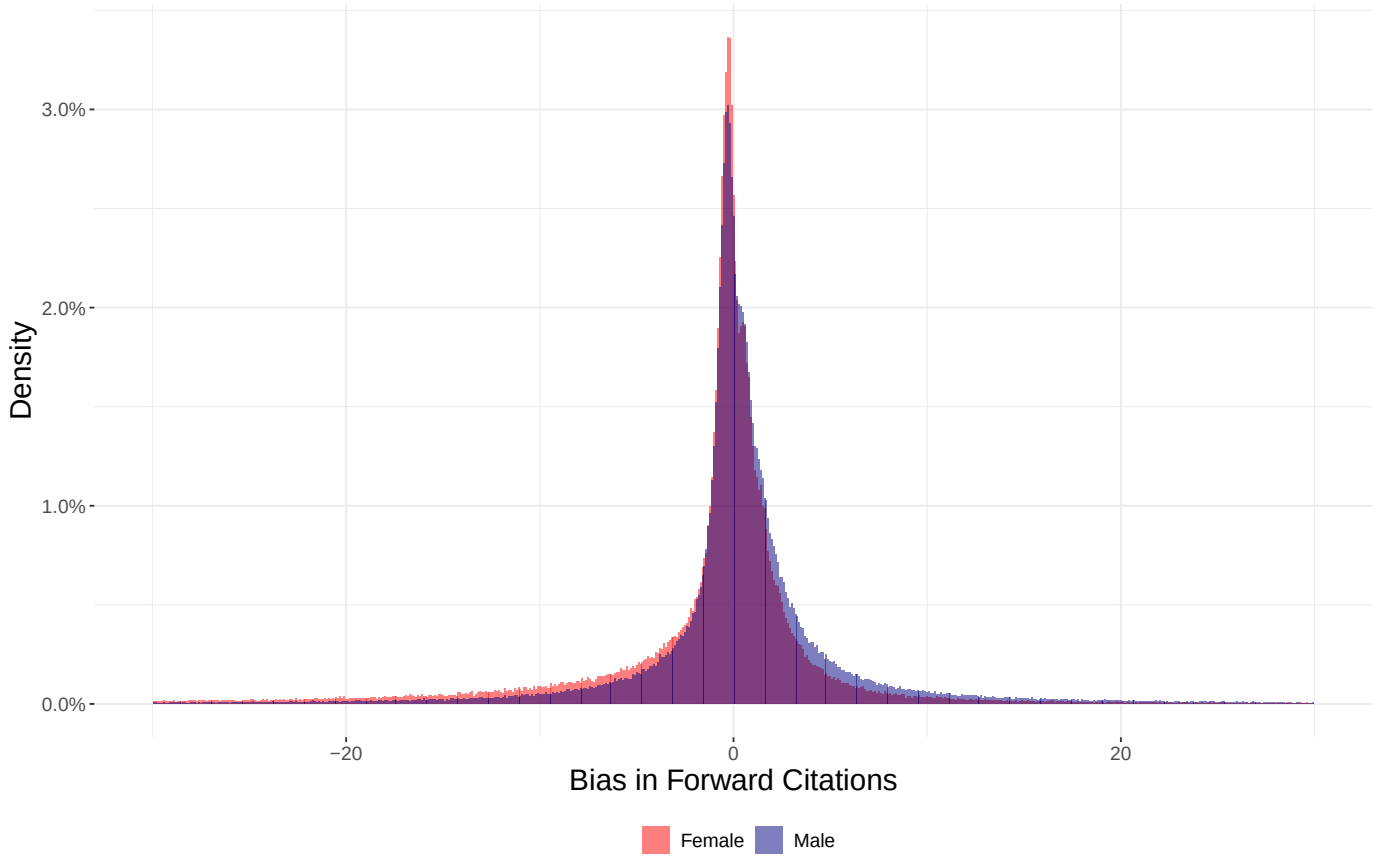


FIGURE 4: UNDERCITATION OF FEMALE LEAD PATENTS

This figure illustrates the difference between forward citations and expected forward citations, as defined by Equation 2. The horizontal axis counts the additional number of citations that a patent should have received after adjusting. Red bars correspond to females, blue bars correspond to males, and purple bars correspond to the overlapping region. The distribution is truncated between -30 to 30.

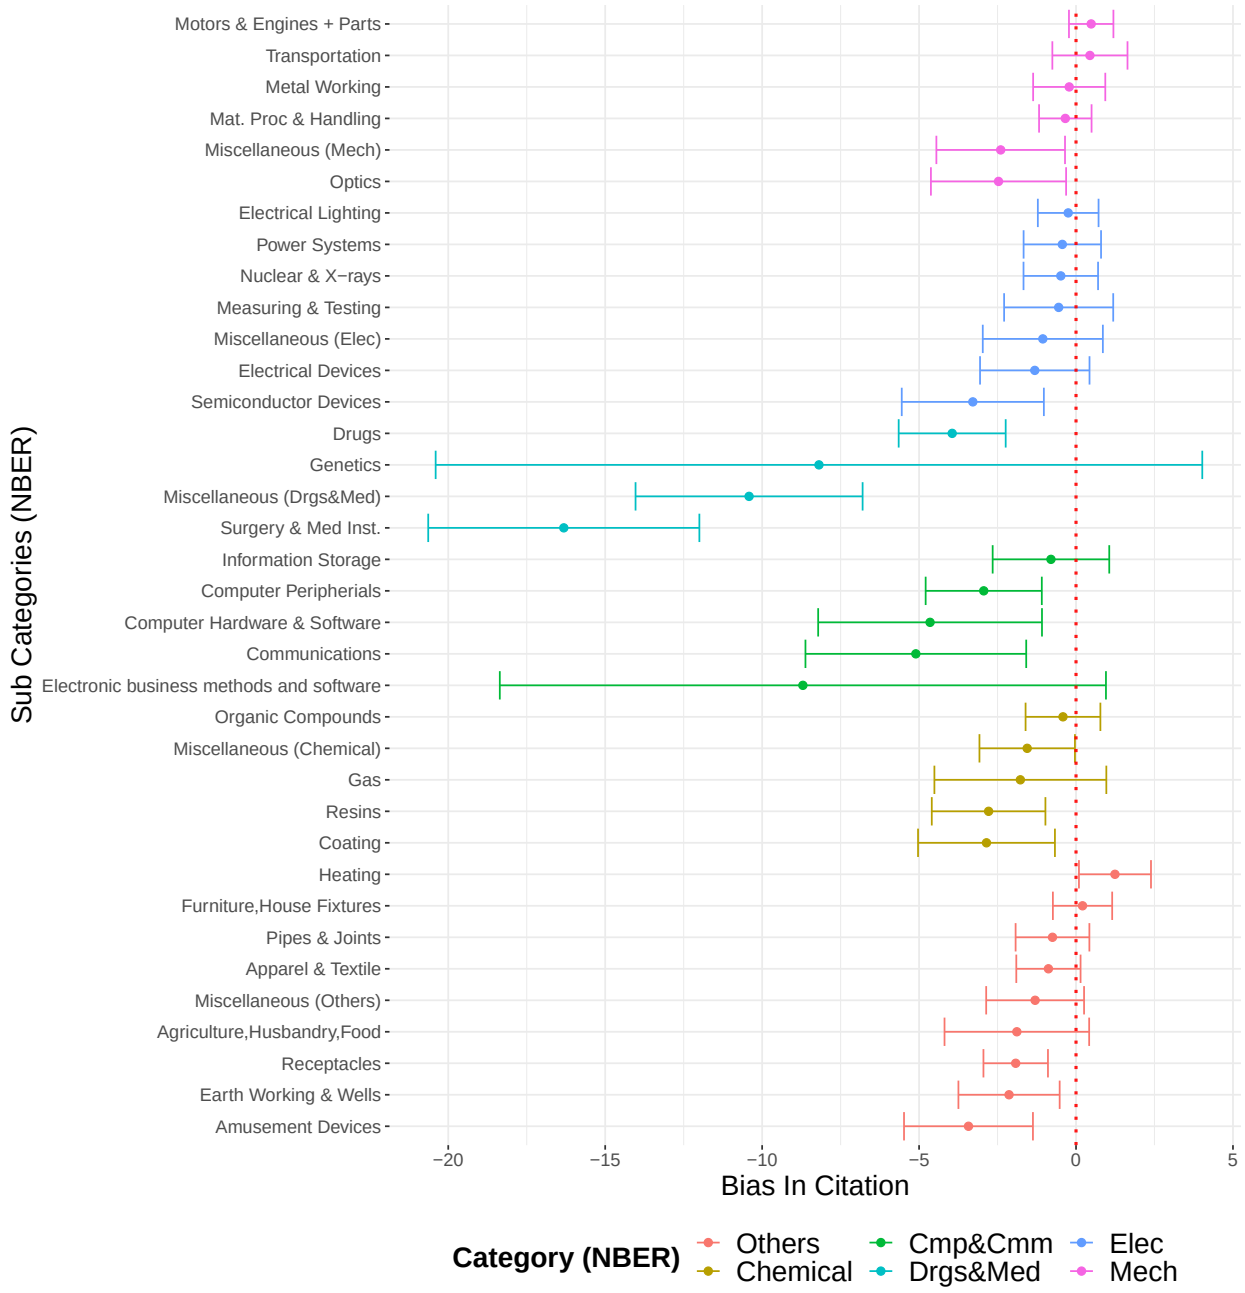


FIGURE 5: DELTA IN CITATIONS BY PATENT SUB CATEGORIES

This figure illustrates the coefficients of Equation 5. For ease of interpretation, each point corresponds to the linear combination of the baseline result for females and the interaction terms, presented in Table IA4. Whiskers correspond to a 95% confidence internal. Coefficients are sorted by by patent category and then by the magnitude of the estimate. Colors corresponds to the patent category as defined by the NBER, where pink observations correspond the mechanical (Mech), purple corresponds to electrical (Elec), blue observations correspond to drugs and medical (Drgs&Med), light green observations correspond to computers and communication (Cmp&Comm), dark green observations correspond to chemical (Chemical), yellow observations correspond to other (Other) categories. The red dotted line is plotted at the zero intercept, representing a no effect.

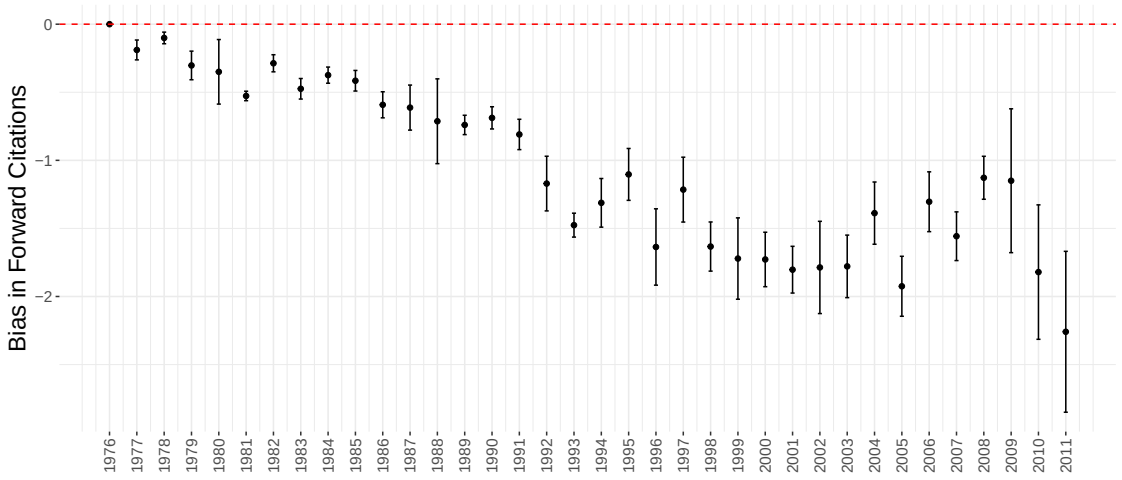


FIGURE 6: EVOLUTION OF DELTA OVER TIME

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.

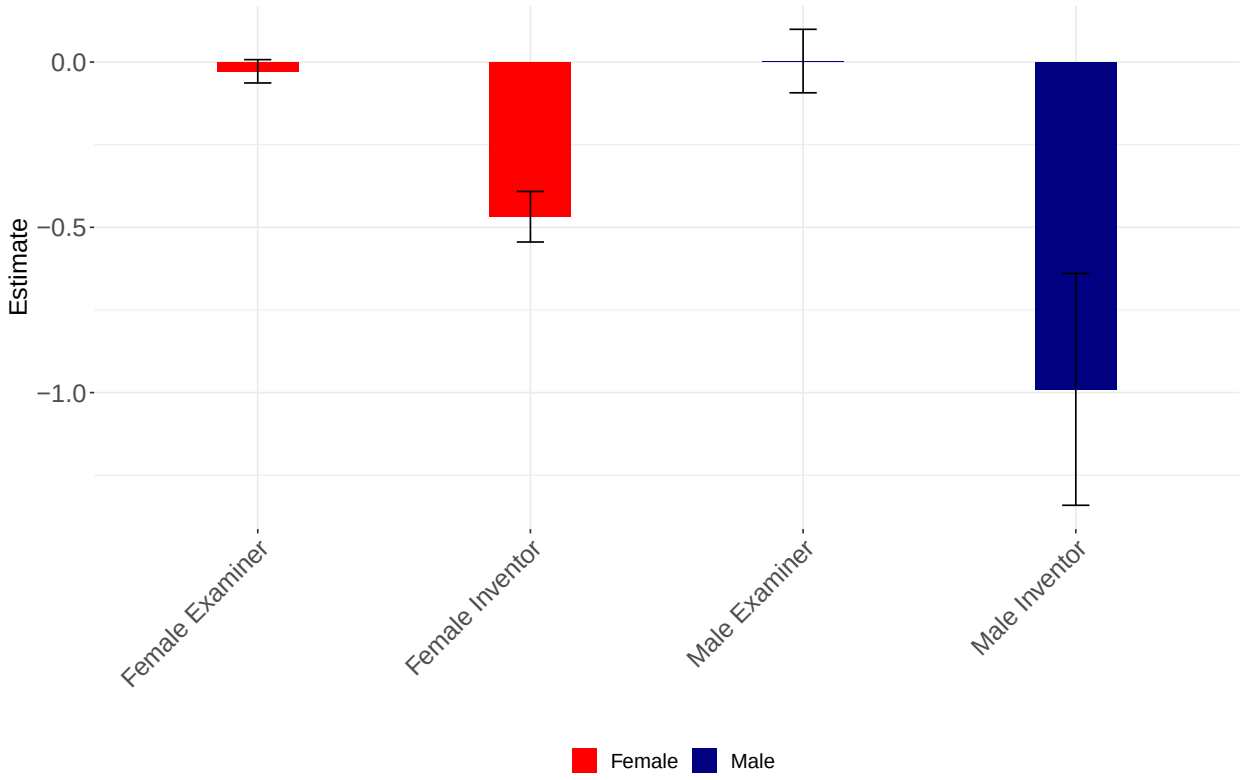


FIGURE 7: DELTA IN EXAMINER AND INVENTOR ADDED CITATIONS

This figure illustrates the delta by examiner and inventor added citations and by gender. The red columns correspond to female-added citations, and the blue columns correspond to male-added citations. Error bars correspond to the 95% confidence interval. Estimates correspond to column (4) of Table 8 and Table 9.

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%, 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	290786	18.09	51.66	236562	15	45.34	F= 518.967***
Panel B: Difference By Top Decile Innovations							
Top Decile	290786			236562			X2= 572.671***
→ No	260016	89%		216170	91%		
→ Yes	30770	11%		20392	9%		
Panel C: Difference by Cooperative Patent Classification							
CPC Section	290733			236533			$\chi^2 = 4871.596^{***}$
→ Chemistry	30780	11%		32379	14%		
→ Electricity	67174	23%		59396	25%		
→ Fixed Constructions	8551	3%		4474	2%		
→ Human Necessities	35268	12%		31598	13%		
→ Mechanical Engineering	24745	9%		13263	6%		
→ Performing Operations	46932	16%		29477	12%		
→ Physics	74182	26%		63730	27%		
→ Textiles	3101	1%		2216	1%		

TABLE 2: FORWARD CITATION WITHOUT MODEL ADJUSTMENTS

This table reports estimates of Equation 1 and studies the number of forward citations by the gender of the lead inventor. The sample includes both patents that received a citation and those that did not receive a citation. 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 customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Extensitve Margin

	Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-3.947^{***}	-1.443^{***}	-1.097^{***}	-0.830^{***}
	(0.266)	(0.400)	(0.280)	(0.270)
Intercept	14.953^{***}			
	(4.497)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	907,996	907,996	907,996	907,996
Adjusted R^2	0.002	0.122	0.199	0.375

Panel B: Intensive Margin

	Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-3.086^{***}	-1.771^{***}	-1.306^{***}	-0.953^{***}
	(0.366)	(0.345)	(0.223)	(0.221)
Intercept	18.091^{***}			
	(2.952)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	527,348	527,348	527,348	527,348
Adjusted R^2	0.001	0.095	0.182	0.354

TABLE 3: MODEL ADJUSTED FORWARD CITATION

This table reports estimates of Equation 2 and studies the number of forward citations by the gender of the lead inventor. The sample includes both patents that received a citation and those that did not receive a citation. 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 customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Extensive Margin

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	−3.796**	−4.057**	−3.952**	−3.263
	(1.547)	(1.537)	(1.571)	(1.962)
Intercept	0.235***			
	(0.085)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	907,996	907,996	907,996	907,996
Adjusted R^2	0.006	0.017	0.051	0.298

Panel B: Intensive Margin

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	−3.791***	−3.853***	−3.769***	−3.312***
	(0.779)	(0.774)	(0.792)	(1.136)
Intercept	1.097***			
	(0.191)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	527,348	527,348	527,348	527,348
Adjusted R^2	0.006	0.011	0.062	0.328

TABLE 4: 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 customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Citations Without Model Adjustment

	Top Decile Patent			
	(1)	(2)	(3)	(4)
Lead Inventor Female	-0.020^{***}	-0.011^{***}	-0.009^{***}	-0.006^{***}
	(0.003)	(0.002)	(0.002)	(0.002)
Intercept	0.106^{***}			
	(0.022)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	527,348	527,348	527,348	527,348
Adjusted R^2	0.001	0.107	0.146	0.185

Panel B: Counterfactual

	Flipped to Top Decile			
	(1)	(2)	(3)	(4)
Lead Inventor Female	0.017^{***}	0.018^{***}	0.018^{***}	0.017^{***}
	(0.004)	(0.004)	(0.004)	(0.006)
Intercept	0.010^{***}			
	(0.002)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	527,348	527,348	527,348	527,348
Adjusted R^2	0.004	0.011	0.016	0.066

TABLE 5: CITATION BY NBER CATEOGRY

This table estimates the difference in citations by NBER Cateogry. Column (1) uses the number of forward citations as its dependent variable, while Column (2) uses the difference in forward citations, as defined by Equation 2. Estimates include interactions for the patent category based on NBER Categories. All specifications include *NBER Cateogry*, *Examiner* × *Customer*, and *Patent Issue Year* fixed effects. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Dependent variable:	
	Forward Citations	Delta in Forward Citations
	(1)	(2)
Female Lead Inventor	−0.077	−3.154***
	(0.395)	(0.337)
Chemical × Female Lead Inventor	−0.602**	−0.512***
	(0.263)	(0.167)
Computers and Communication × Female Lead Inventor	−1.320***	−2.837***
	(0.270)	(0.990)
Drugs and Medical × Female Lead Inventor	−5.797***	−7.260***
	(0.345)	(0.638)
Electrical × Female Lead Inventor	−0.767**	0.025
	(0.347)	(0.147)
Mechanical × Female Lead Inventor	0.096	0.572***
	(0.263)	(0.157)
NBER Category FE	Yes	Yes
Examiner x Customer FE	Yes	Yes
Examiner Art Unit FE	Yes	Yes
Patent Issue Year FE	Yes	Yes
Observations	436,458	436,458
Adjusted R^2	0.296	0.282

TABLE 6: FORWARD CITATIONS, EMERGING FIELDS

This table studies the citations to new fields of innovation. *New Field* takes the value of one if the art unit first appeared within five years of the patent being granted. The dependent variable is the difference in the observed number and the expected number of citations for a patent, as defined by Equation 2. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Forward Citations

	Forward Citations			
	(1)	(2)	(3)	(4)
Emerging Field	−0.099	0.032	−0.071	−0.541
	(2.506)	(0.551)	(0.534)	(0.344)
Female Lead Inventor	−3.351***	−1.703***	−1.324***	−1.013***
	(0.524)	(0.339)	(0.220)	(0.160)
Emerging Field × Female Lead Inventor	0.950	−0.246	0.065	0.282
	(0.780)	(0.430)	(0.324)	(0.374)
Intercept	18.119***			
	(3.486)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	527,348	527,348	527,348	527,348
Adjusted R^2	0.001	0.095	0.182	0.354

Panel B: Delta in Forward Citations

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Emerging Field	−0.016	0.213	0.308	0.084
	(0.185)	(0.352)	(0.312)	(0.202)
Lead Female Inventor	−3.612***	−3.693***	−3.637***	−3.305***
	(0.823)	(0.831)	(0.847)	(1.141)
Emerging Field × Lead Female Inventor	−0.643	−0.573	−0.484	−0.036
	(0.603)	(0.598)	(0.568)	(0.298)
Intercept	1.102***			
	(0.183)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	527,348	527,348	527,348	527,348
Adjusted R^2	0.006	0.011	0.062	0.328

TABLE 7: YEARS AFTER PATENT IS GRANTED

This table estimates Equation 1 and studies the difference in forward citations by the number of years after the patent was granted. 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*, *Examiner* $\times$ *Customer*, and *Patent Grant Year* fixed effects. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations			
	0-1 Years	2-5 Years	6-10 Years	11-20 Years
	(1)	(2)	(3)	(4)
Lead Inventor Female	-0.369^{***}	-0.626^{***}	-1.185^{***}	-2.479^{***}
	(0.105)	(0.066)	(0.108)	(0.252)
Examiner x Customer FE	Yes	Yes	Yes	Yes
Art Unit FE	Yes	Yes	Yes	Yes
Patent Issue Year FE	Yes	Yes	Yes	Yes
Observations	12,238	237,734	258,853	223,519
Adjusted R^2	-1.241	-0.247	0.463	0.467

TABLE 8: EXAMINER-ADDED CITATIONS

This table studies the source of examiner-added citations for male inventors. The dependent variable is the difference in forward citations. 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 customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Citation Added by Female Lead Examiner

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-0.029^{***}	-0.021^{**}	-0.014	-0.028
	(0.011)	(0.009)	(0.011)	(0.018)
Intercept	0.131^{***}			
	(0.038)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	66,757	66,757	66,757	66,757
Adjusted R^2	0.0001	0.017	0.062	0.101

Panel B: Citation Added by Male Lead Examiners

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-0.006	0.009	0.004	0.003
	(0.034)	(0.038)	(0.037)	(0.049)
Intercept	0.060			
	(0.065)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	180,397	180,397	180,397	180,397
Adjusted R^2	-0.00000	0.010	0.046	0.219

TABLE 9: INVENTOR-ADDED CITATIONS

This table studies the source of inventor-added citations. The dependent variable is the difference in forward citations. 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 customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Citation Added by Female Lead Inventors

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-0.508^{***}	-0.461^{***}	-0.473^{***}	-0.468^{***}
	(0.051)	(0.053)	(0.033)	(0.039)
Intercept	0.609^{***}			
	(0.063)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	29,984	29,984	29,984	29,984
Adjusted R^2	0.004	0.019	0.075	0.223

Panel B: Citation Added by Male Lead Inventors

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-1.363^{***}	-1.360^{***}	-1.306^{***}	-0.990^{***}
	(0.174)	(0.174)	(0.172)	(0.178)
Intercept	1.083^{***}			
	(0.171)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	282,370	282,370	282,370	282,370
Adjusted R^2	0.001	0.010	0.091	0.191

TABLE 10: 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 customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	$\log(dollar)$			
	(1)	(2)	(3)	(4)
Female Lead Inventor	0.083***	−0.014	−0.020	−0.050
	(0.029)	(0.033)	(0.029)	(0.041)
Forward Citation	0.00003	−0.0001	0.0001	0.0001
	(0.0003)	(0.0003)	(0.0002)	(0.0003)
$\widehat{\text{Forward Citation}}$	0.004***	0.004***	0.002***	0.003***
	(0.001)	(0.0003)	(0.001)	(0.001)
Intercept	0.716***			
	(0.138)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	202,865	202,865	202,865	202,865
Adjusted R^2	0.008	0.127	0.438	0.321

INTERNET APPENDIX
FOR ONLINE PUBLICATION

Appendix A Explanation of Causal BERT (C-BERT)

C-BERT 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-BERT to identify the effect of gender on the impact of patents, we first need to train the model. As shown in Figure IA2, 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: a BERT model for generating text embeddings, a logit-linear model the 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-BERT 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 BERT 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 are 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 corresponds to one patent.

Appendix A Encoder Architecture

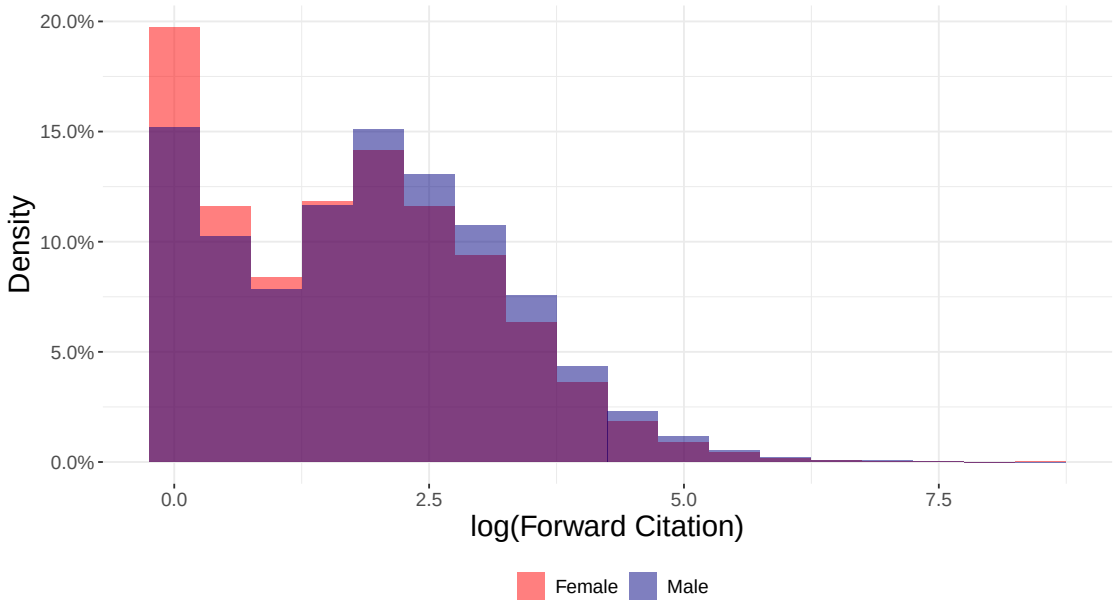
The encoder architecture works as follows. Let W denote the original input sentence in words. As shown in Figure IA3, before entering the encoder, W is broken down into three parts: a token embedding E_W^T, which represents the content of the sentence; a segmentation embedding E_W^S, which labels tokens 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 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 on to 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 tokens 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.

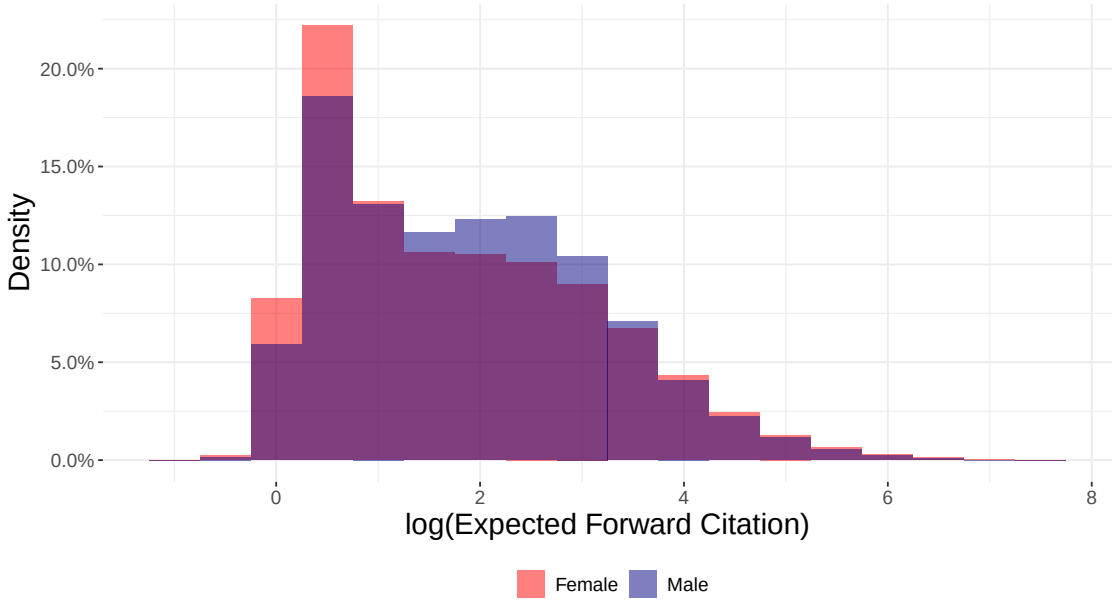
The output of this multi-headed attention layer is then passed through a normalization layer with residual connection. Residual connection is achieved by passing the input of the multi-headed attention layer directly to the normalization layer along with the output of the multi-headed attention layer. 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 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 BERT model uses the encoder architecture to train for two tasks:

masked language modeling (MLM) and next sentence prediction (NSP). 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. To train for next sentence prediction, the encoder takes pairs of sentences as inputs and predicts whether the second sentence should appear after the first sentence. This loss is computed using binary-cross-entropy. The final trained BERT model can output embeddings of sentences or entire texts that represents not only the meaning of the tokens but also the contextual relationship between tokens and sentences.



(a) Forward Citations



(b) Model Implied Forward

FIGURE IA1: 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. Red bars correspond to females, blue bars correspond to males, and purple bars correspond to the overlapping region.

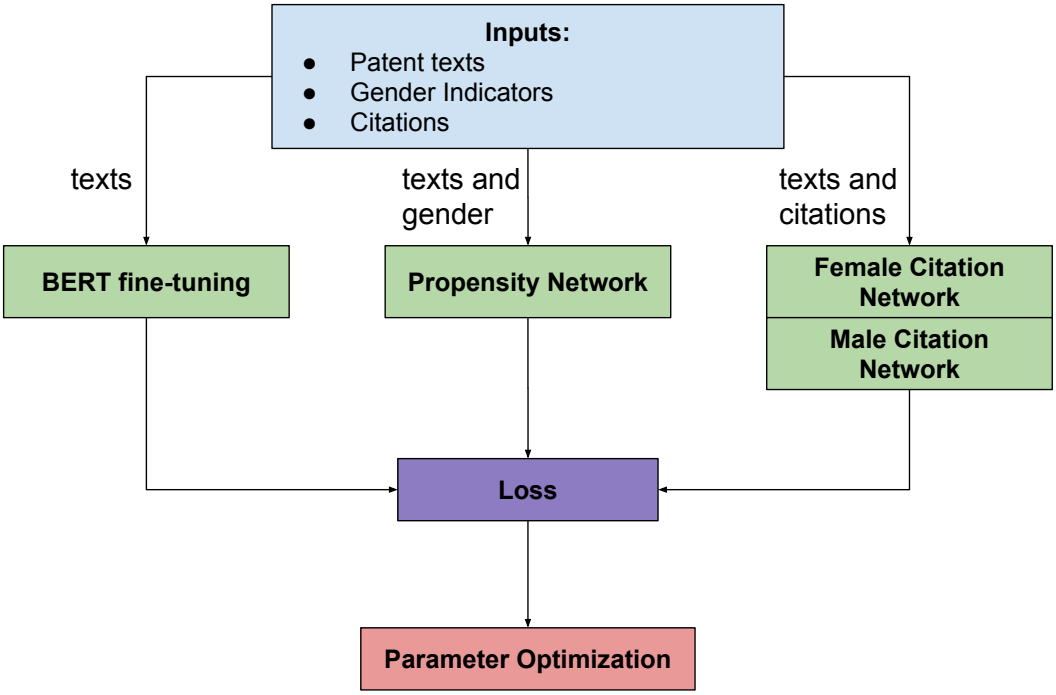


FIGURE IA2: C-BERT TRAINING PROCEDURE

The figure illustrates the training procedure of C-BERT 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.

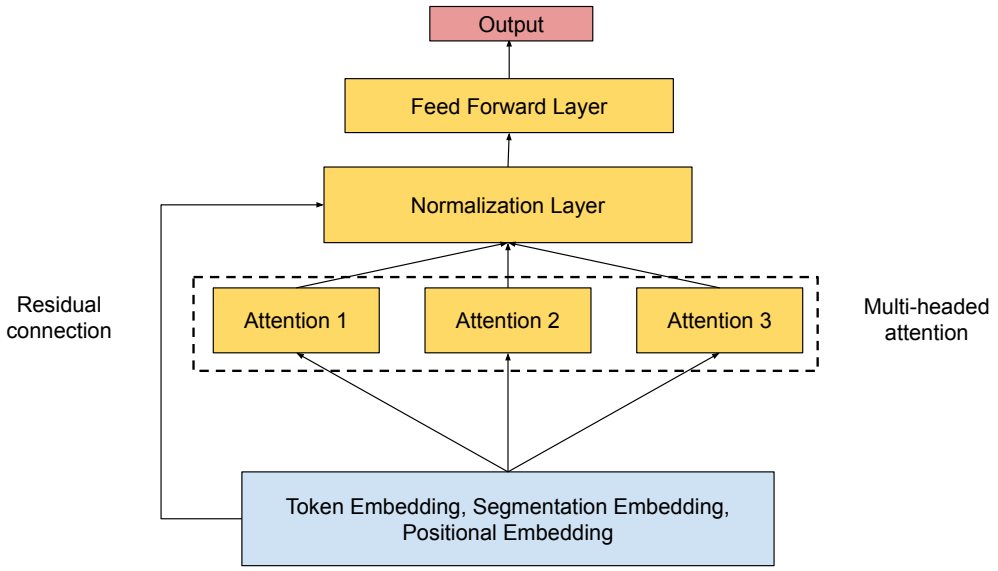


FIGURE IA3: ENCODER MODULE

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.

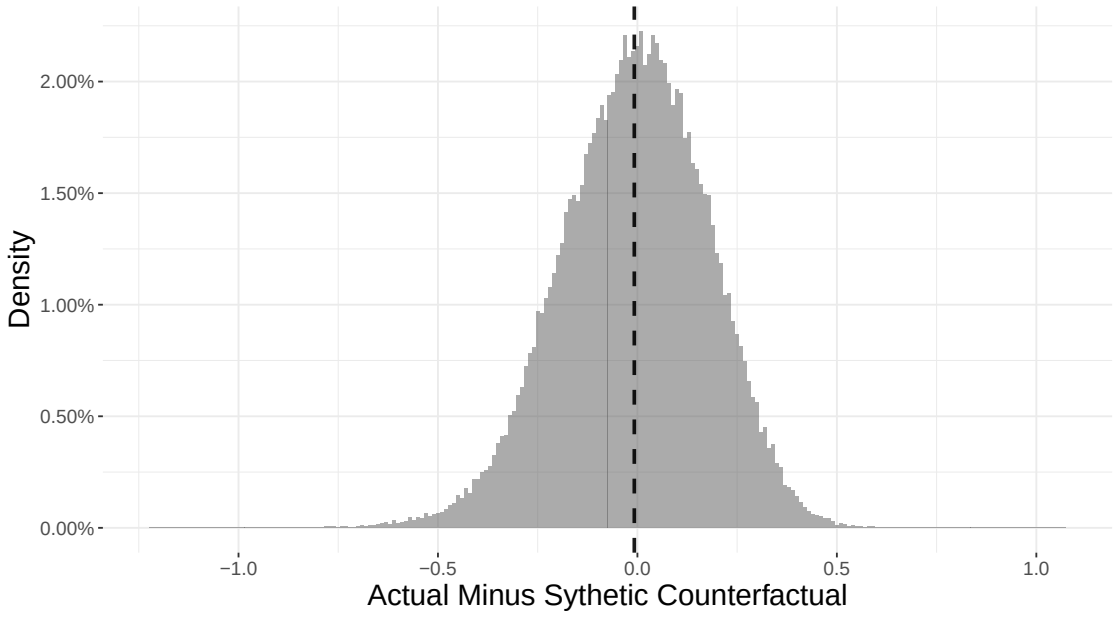


FIGURE IA4: SYNTHETIC DATA TESTS

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).

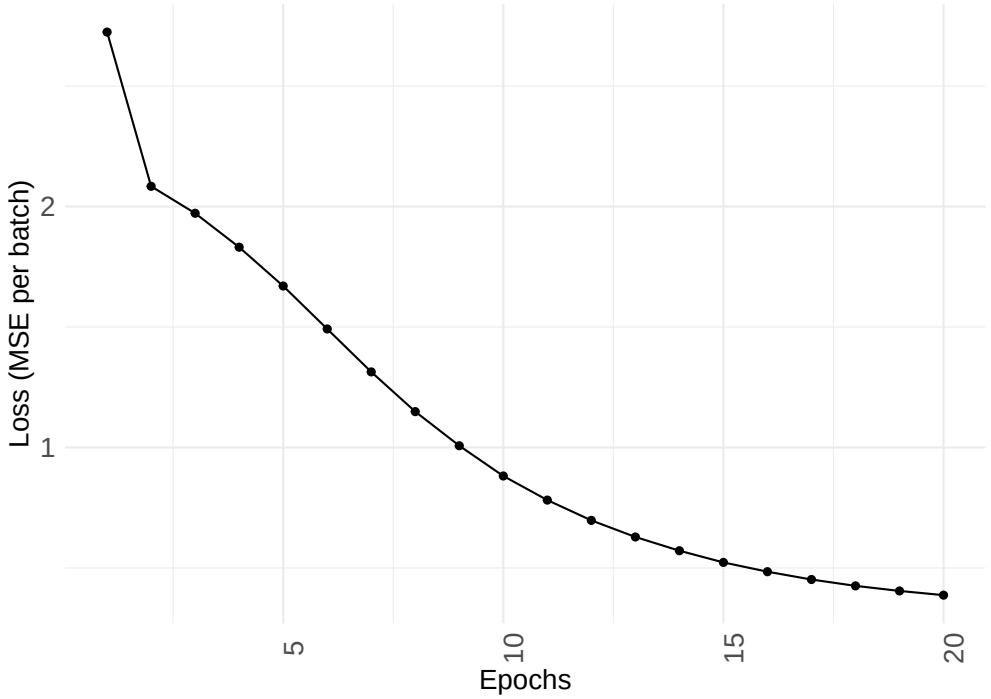


FIGURE IA5: LOSS FUNCTION

This figure illustrates the loss function of the C-BERT 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 IA1: DIFFERENCE IN WRITING STYLES

This table reports the difference in the writing style between males and females. The sample covers patents issued from 1976-01-01 through 2021-12-31 and has at least 120 words. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO, Google Patents.

Gender	Male			Female			
	N	Mean	SD	N	Mean	SD	Test
Number of Words	1981500	159.64	43.26	136152	155.22	38.51	F= 1346.691***
Sentiment Score	1981500	0.04	0.13	136152	0.05	0.13	F= 35.241***
Sentiment	1981500	0.36	0.9	136152	0.37	0.9	F= 6.967***
Flesch-Kincaid	197487	23.33	15.01	14278	23.87	15.12	F= 17.13***
Flesh	197487	9.78	41.76	14278	7.04	42.2	F= 57.028***
Gunning-Fog	197487	27.26	15.91	14278	27.66	16.04	F= 8.487***
Coleman-Liau	197487	13.4	2.44	14278	13.63	2.61	F= 119.976***
Dale-Chall	197487	11.88	2.42	14278	12.17	2.42	F= 194.12***
Ari	197487	26.64	19.66	14278	27.17	19.81	F= 9.695***
Linsear-Write	197487	34.08	28.03	14278	34.78	28.54	F= 8.082***
Spache	197487	11.57	5.58	14278	11.75	5.62	F= 14.111***

TABLE IA2: ASSOCIATION BETWEEN GENDER AND FORWARD CITATION

This table uses all patent observations without applying C-BERT. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-3.621^{***}	-1.023^{***}	-0.784^{***}	-0.591^{**}
	(0.276)	(0.310)	(0.240)	(0.256)
Intercept	15.185^{***}			
	(4.520)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	6,312,796	6,312,796	6,312,796	6,312,796
Adjusted R^2	0.0004	0.126	0.206	0.318

TABLE IA3: 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 the difference in forward citations, as defined by Equation 2. Estimates include interactions for the patent category based on CPC Section. All specifications include *CPC Section*, *Examiner* × *Customer*, and *Patent Issue Year* fixed effects. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Dependent variable:	
	Forward Citations	Delta in Forward Citations
	(1)	(2)
Lead Female Inventor	-1.192^{*}	-3.537^{***}
	(0.685)	(0.821)
Electricity × Lead Female Inventor	0.054	0.643
	(0.568)	(0.472)
Fixed Constructions × Lead Female Inventor	0.280	1.676^{***}
	(0.507)	(0.229)
Human Necessities × Lead Female Inventor	-1.476^{***}	-3.286^{***}
	(0.350)	(1.007)
Mechanical Engineering × Lead Female Inventor	1.356^{*}	2.603^{***}
	(0.768)	(0.649)
Performing Operations × Lead Female Inventor	0.796^{*}	1.037^{***}
	(0.400)	(0.092)
Physics × Lead Female Inventor	0.435	0.444
	(0.453)	(0.597)
Textiles × Lead Female Inventor	2.429^{**}	1.021^{**}
	(0.998)	(0.487)
CPC Section FE	Yes	Yes
Examiner x Customer FE	Yes	Yes
Examiner Art Unit FE	Yes	Yes
Patent Issue Year FE	Yes	Yes
Observations	527,266	527,266
Adjusted R^2	0.356	0.329

TABLE IA4: CITATION BY PATENT NBER SUB CATEGORY

This table estimates the difference in citations by NBER subcategories. Column (1) uses the actual number of citations as its dependent variable, while Column (2) uses the *Delta* in citations, as defined by Equation 2. Estimates include interactions for the patent subcategory based on NBER classifications. All specifications include *Patent Subcategory*, *Examiner* × *Customer*, and *Patent Issue Year* fixed effects. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Dependent variable:	
	Forward Citations	Delta in Forward Citations
	(1)	(2)
Female Lead Inventor	0.499	−1.915***
	(0.470)	(0.440)
Agriculture,Husbandry,Food× Female Lead Inventor	0.204	−1.885
	(0.674)	(1.176)
Amusement Devices× Female Lead Inventor	1.488	−3.426***
	(2.544)	(1.048)
Apparel & Textile× Female Lead Inventor	2.114***	−0.879
	(0.662)	(0.522)
Coating× Female Lead Inventor	−1.591	−2.853**
	(0.967)	(1.113)
Communications× Female Lead Inventor	−2.118***	−5.103***
	(0.528)	(1.794)
Computer Hardware & Software× Female Lead Inventor	−1.807**	−4.650**
	(0.883)	(1.819)
Computer Peripherials× Female Lead Inventor	0.546	−2.940***
	(0.955)	(0.944)
Drugs× Female Lead Inventor	−4.586***	−3.945***
	(1.045)	(0.870)
Earth Working & Wells× Female Lead Inventor	−2.061***	−2.134**
	(0.637)	(0.823)
Electrical Devices× Female Lead Inventor	−1.157**	−1.314
	(0.473)	(0.890)
Electrical Lighting× Female Lead Inventor	−2.045*	−0.249
	(1.195)	(0.494)
Electronic business methods and software× Female Lead Inventor	−8.866**	−8.702*
	(3.623)	(4.926)
Furniture,House Fixtures× Female Lead Inventor	−0.214	0.207
	(0.792)	(0.482)
Gas× Female Lead Inventor	0.102	−1.773
	(1.447)	(1.397)
Genetics× Female Lead Inventor	−8.841*	−8.189
	(4.996)	(6.230)
Heating× Female Lead Inventor	−1.353	1.241**
	(1.130)	(0.585)
Information Storage× Female Lead Inventor	−2.166**	−0.798
	(0.879)	(0.948)

continued on next page...

TABLE IA4: CITATION BY PATENT NBER SUB CATEGORY (CONTINUED)

	Dependent variable:	
	Forward Citations	Delta in Forward Citations
	(1)	(2)
Mat. Proc & Handling× Female Lead Inventor	0.259	−0.341
	(0.665)	(0.427)
Measuring & Testing× Female Lead Inventor	−0.728	−0.553
	(0.513)	(0.887)
Metal Working× Female Lead Inventor	0.509	−0.217
	(1.003)	(0.586)
Miscellaneous (Chemical)× Female Lead Inventor	−1.127*	−1.556*
	(0.582)	(0.776)
Miscellaneous (Drgs&Med)× Female Lead Inventor	−5.059	−10.416***
	(3.144)	(1.845)
Miscellaneous (Elec)× Female Lead Inventor	−0.584	−1.058
	(0.592)	(0.976)
Miscellaneous (Mech)× Female Lead Inventor	−1.440	−2.400**
	(1.775)	(1.045)
Miscellaneous (Others)× Female Lead Inventor	−0.532	−1.303
	(0.927)	(0.795)
Motors & Engines + Parts× Female Lead Inventor	0.0002	0.483
	(0.662)	(0.361)
Nuclear & X-rays× Female Lead Inventor	−2.954***	−0.485
	(0.961)	(0.606)
Optics× Female Lead Inventor	−2.472	−2.469**
	(1.711)	(1.099)
Organic Compounds× Female Lead Inventor	−1.812	−0.415
	(1.354)	(0.608)
Pipes & Joints× Female Lead Inventor	−4.125**	−0.750
	(1.746)	(0.600)
Power Systems× Female Lead Inventor	−0.896	−0.436
	(0.565)	(0.630)
Receptacles× Female Lead Inventor	−2.042	−1.921***
	(1.275)	(0.524)
Resins× Female Lead Inventor	−1.825**	−2.785***
	(0.824)	(0.924)
Semiconductor Devices× Female Lead Inventor	−2.144**	−3.288***
	(0.891)	(1.155)
Surgery & Med Inst.× Female Lead Inventor	−7.317***	−16.319***
	(1.159)	(2.204)
Transportation× Female Lead Inventor	−0.866*	0.443
	(0.438)	(0.611)
Patent Subcategory (NBER) FE	Yes	Yes
Examiner x Customer FE	Yes	Yes
Examiner Art Unit FE	Yes	Yes
Patent Issue Year FE	Yes	Yes
Observations	436,458	436,458
Adjusted R^2	0.299	0.283

TABLE IA5: ROBUSTNESS TO SAMPLE SELECTION

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

Panel A: Single Author

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-3.431^{***}	-3.376^{***}	-3.234^{***}	-2.764^{***}
	(0.614)	(0.598)	(0.628)	(0.823)
Intercept	1.456^{***}			
	(0.261)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	124,280	124,280	124,280	124,280
Adjusted R^2	0.002	0.015	0.138	0.421

Panel B: Inventors Same Gender

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-4.042^{***}	-3.984^{***}	-3.931^{***}	-3.769^{***}
	(0.338)	(0.323)	(0.296)	(0.163)
Intercept	1.725^{***}			
	(0.239)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	129,298	129,298	129,298	129,298
Adjusted R^2	0.003	0.017	0.211	0.466

TABLE IA6: BIAS IN FORWARD CITATION, USING LONGFORMER

This table replaces the BERT model using Longformer to study the full text of patents. The sample covers patents ranging from 4800000 through 5299999. These range roughly from 1986 until 1994. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	−1.961***	−2.108***	−2.198***	−1.993***
	(0.703)	(0.548)	(0.220)	(0.201)
Intercept	4.767***			
	(0.471)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	37,020	37,020	37,020	37,020
Adjusted R^2	0.001	0.020	0.255	0.292

TABLE IA7: PLACEBO, RANDOM SAMPLE

This table presents a placebo test by randomizing the gender of patents and re-running our C-BERT approach to establish the effects are not an artifact of C-BERT. Panel A uses the number of forward citations as the dependent variable, while Panel B uses bias computed using C-BERT. The sample covers patents issued from 1976-01-01 through 2021-12-31. Standard errors are clustered at the patent customer and patent issue year level. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively. Data source: USPTO.

Panel A: Forward Citations

	Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	0.022	−0.027	−0.066	−0.106
	(0.154)	(0.145)	(0.106)	(0.096)
Intercept	19.918***			
	(3.146)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	471,461	471,461	471,461	471,461
Adjusted R^2	−0.00000	0.109	0.202	0.304

Panel B: Delta in Forward Citations

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	−0.042	−0.036	−0.058	−0.214
	(0.202)	(0.208)	(0.212)	(0.250)
Intercept	−3.626***			
	(0.537)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	471,461	471,461	471,461	471,461
Adjusted R^2	−0.00000	0.025	0.051	0.221

TABLE IA8: BIAS IN FORWARD CITATION, USING SCIBERT

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

	Delta in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-7.524^{***}	-7.693^{***}	-7.584^{***}	-6.580^{***}
	(1.536)	(1.458)	(1.514)	(2.115)
Intercept	0.680^{***}			
	(0.141)			
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Art Unit FE	No	Yes	Yes	Yes
Patent Issue Year FE	No	Yes	Yes	Yes
Observations	602,974	602,974	602,974	602,974
Adjusted R^2	0.012	0.033	0.064	0.195