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

Yael V. Hochberg[‡] **Ali Kakhbod**[§] **Peiyao Li**[||] **Kunal Sachdeva**[¶]

January 2023

Abstract

Implementing a state-of-the-art machine learning technique for causal identification from text data, we document that women are under-cited relative to the quality of their patents. For the equivalent patent with a lead female inventor, a patent with a male lead inventor would have received 27% more citations. These effects are magnified for corporate innovations. 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. The under-recognition of female-authored patents likely has implications for the allocation of talent in the economy.

JEL classification: J16, J71, O30

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, Asaf Manela, Danqing Mei, S. Lakshmi Naaraayanan, Vira Semenova, Amit Seru, and Carolyn Stein for their helpful comments and suggestions.

[‡]Jesse H. Jones Graduate School of Business, Rice University, hochberg@rice.edu

[§]Haas School of Business, University of California at Berkeley, akakhbod@berkeley.edu

[||]Haas School of Business, University of California at Berkeley, ojhfklsjhl@berkeley.edu

[¶]Jesse H. Jones Graduate School of Business, Rice University, kunal.sachdeva@rice.edu

Female inventors appear to face significant obstacles when seeking patents. Women face disparities in the approval of patents (Gavrilova and Juranek (2021); Jensen, Kovács, and Sorenson (2018)), and they are underrepresented among inventors in patent applications more generally (Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019); Subramani, Aneja, and Reshef (2021)). A possible corollary of these findings is a selection effect: assuming a patent is applied for and granted, patents with female inventors may be of higher quality than average and, as a result, better cited. Surprisingly, however, standard analysis of patent citation suggests that female-authored patents receive fewer citation relative to male-authored patents. This difference in citation patterns would suggest that either female inventors produce lower quality patents, or, more concernedly, their patent quality is 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, however, that mimic this ideal or suitable instrumental variables have been elusive.

We utilize a novel machine learning technique that allows measurement of the causal contribution of gender to citation of patents of similar quality—causal bidirectional encoder representation from transformers or C-BERT (Veitch, Sridhar, and Blei (2020)).[1] C-BERT estimates causal effects from observational text data, adjusting for confounding features of the text, such as the subject or writing quality. It assumes that the text content suffices for causal identification but is prohibitively complex for standard anal-

[1]This approach builds on leading approaches proposed in the computer science literature, including recent papers by Khetan et al. (2022); Shao et al. (2021).

ysis. 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).

The intuition behind our use of C-BERT 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. 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. Figure 1 illustrates the assumption of the C-BERT mediating approach. The next step is to train the model. First, we use a pre-trained BERT model provided by Google (TensorFlow) to transform the text of each patent 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 using a propensity score estimation network that matches on text across gender. 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 females are the lead authors of the patents, while the second mapping is trained using the data where males are the lead authors. 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. The procedure is illustrated in Figure 2.

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 label each patent by the inferred gender of the inventors following Desai (2019), textually matching the first names of the lead inventor to data from the Social Security Administration and the world intellectual patent organization.

Our analysis begins by documenting that, even when no adjustments are made, there is a statistically significant difference in the number of forward citations for patents authored by women versus men in our matched data. This pattern persists even 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 patent, and the patent issue year. The distributional plots of forward citations confirm the statistical tests and suggest that citations of female-authored patents are slightly skewed downward.

After we controlled for differences in the quality of patents based on the gender of the lead author using C-BERT, we find that women received approximately 27% fewer citations than men for equivalent patents in the same art unit and evaluated by the same examiner. This equates to approximately 4.6 fewer citations per patent. The impact of this bias is pronounced for the most highly cited patents, with female-authored patents being roughly 20% less likely to reach the top decile of citations. These results suggest that, while patents lead-authored by women would receive additional citations if they were lead-authored by men, but the average treatment effect of being a female mask these differences. This is consistent with the possibility that female-authored patents are only approved if they are of higher quality than equivalent male-authored patents.

Our results hold across NBER's major categories and subcategories of patent technology, with some heterogeneity in subcategories.[2] The effects are ameliorated in emerging technology fields, where patents with female lead inventors do not face under citation, hinting at the possibility that females do not face established barriers in these newer fields. The results are robust to a variety of alternative specifications and do not seem to be attributable to sample selection, our definition of gender, or model overfitting.

We next examined the role of inventors and examiners in the undercitation of patents with female lead inventors. Our analysis found that, compared to patents with female lead inventors, patents with male lead inventors undercite past patents written by female lead inventors more. In contrast, both male and female examiners appear to slightly undercite female-led patents. These examiner effects are relatively small in terms of their economic impact, compared to the overall undercitation effect. When controlling for art units and other patent characteristics, these results suggest that the undercitation of female patents is largely due to male lead inventors.

There are a few limitations to our analysis that should be noted. First, C-BERT relies on the assumption that the text of the patent captures all the factors that influence the number of forward citations. While it is not possible to test this assumption directly, patents may be an ideal setting for this approach because the text of the patent is closely related to the outcome of interest. We also included controls for potential confounding factors that are not accounted for by our neural networks. Another potential concern is the accuracy of our model in computing counterfactual outcomes. To address this, we conducted a semi-synthetic dataset exercise to demonstrate the high accuracy of our model. A third concern is the possibility of overfitting, which is a common issue with machine-learning approaches. We mitigated this risk by reducing the number of iterations of the entire training dataset (EPOCH) and found similar results.

Our results of course should be interpreted with care. Although our evidence sug-

[2]We also consider this test using the Cooperative Patent Classification (CPC).

gests 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.

That said, our findings have potentially important implications. First, the literature has highlighted that innovation is motivated by the expected profits derived from the property rights granted to patentees, Moser (2005, 2013).[3] If women are not equally recognized for equivalent patents, this may discourage them from entering the innovation economy, potentially reducing contributions from half of the population. Too, this may exacerbate the already substantial wedge between men and women in science, technology, engineering and mathematics (STEM) fields (Beede, Julian, Langdon, McKittrick, Khan, and Doms, 2011), leading to further inefficient allocation of labor.

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 paper makes a number of additional contributions to the literature. First, 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 identify-

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

ing differences in patent applications and approvals, our findings focus on a relatively unexplored question: whether women also face obstacles in citation of their patents.

Second, we 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.

Third, 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.

Finally, given the challenges of replicating results in empirical research, we intend to make our code and data fully available via a GitHub repository for researchers as a public good. This repository will allow future researchers to apply C-BERT to other applications in finance and economics while also allowing for simple replication and extensions to all results presented here.

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. 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. Expanding our analysis to the full text would render the analysis intractable.

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 demonstrate that this measure is systematically biased down for female-authored patents. We therefore distinguish between patent citations (the 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.

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

1.3 Gender of Inventors and Examiners

Our main treatment variable is the gender of the lead inventor (first author). One data limitation with this, however, is that inventors do not disclose their gender when applying for a patent. Because of this, we must infer them 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%.[4] 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., 2012). 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 construct a subsample of patents with male lead inventors. We then estimate a propensity model using a one layer logit-linear neural network, where the objective function is the binary-cross-entropy between the predicted treatment indicator and the true treatment indicator. The output of this neural network is the probability that this patent is written by a female lead author.

Next we randomly select a subsample of male authored patents from the full dataset (without replacement). 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%). The intuition is that, if there is no matching sample in the male data that is close in nature to a patent written by a female lead-author, the patent is dropped, and, if there is no matching sample in the female data that is close in nature to a patent written by a male lead-author, the patent is also dropped.

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, the inventors cite precedent

[4]We take a conservative approach and apply a high confidence interval to reduce Type I errors when identifying males and females. We discuss this further in the data and robustness section.

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 started to disclose 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 samples of patents from 1976 through 2021. Note, these samples only 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 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. 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 digestible manner. In further robustness tests we also consider the Cooperative Patent Classification (CPC) of the patents.

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.

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 ultimate patent forward 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.[5] 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.

The encoder architecture works as follows. Let W denote the original input sentence in words. As shown in Figure 3, 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

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

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.

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 Represen-

tations from Transformers (*C-BERT*) that allows us to casually estimate the contribution 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). In this paper, we apply this new technique 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 IA1, 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 two 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. As mentioned above, 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. 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. In this section, we explain our assumptions, assess their validity, and provide a

checklist for future researchers 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 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. Studying the quality of our C-BERT model, we find that it discovers the 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, we follow the procedure of Veitch et al. (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? To answer this, we first demonstrate that there appears on the surface to be a modest difference in the citation counts for female- and male-authored patents when using a standard data analysis approach. This simple analysis, however, masks the underlying true causal effects of gender on patent citation. We next use C-BERT to mediate for the quality of the patent, and show that women lead inventors receive fewer citations to their patents than they would have received if they were male.

3.1 Comparing Between Genders Without Model Adjustments

Plotting the unconditional differences in citations by gender, we can illustrate the similarity in patent citations. The histogram for citations for male and female lead inventors, plotted in Panel A of Figure 4, visually demonstrates female receive fewer citations than males.[6] Estimating the mean difference in citations between genders suggests that male lead inventors receive statistically more citations than female lead inventors, on average (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 citations come from a different distribution.

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, forward citations. Our specification includes fixed effects for customer-examiner pair ($\delta_{customer \times examiner}$), art unit ($\delta_{ArtUnit}$), and year of grant ($\delta_{GrantYear}$), to control for time-invariant heterogeneity across patent applications as well as time trends. All 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. The estimates from this table suggests that female lead investors receive between 1.2 to 2.9 fewer citations. The fact that we find a significant undercitation of female lead investors could be considered surprising, given the obstacles female's face when patent filing and granting, as documented elsewhere.

[6] To address skewness and show more clearly, we the natural logarithm of citations and present a histogram in Panel A of Figure IA2

That literature would also suggest a selection effect, whereby patents granted to female inventors should be of higher quality and thus, using forward citations as the accepted measure of patent quality, should receive more citations. This expectation is supported by the literature. For example, in economics, Card, DellaVigna, Funk, and Iriberri (2020) reports that female-authored papers on average receive 25% more citations than similar male-authored papers. Similarly, Hengel and Moon (2020) report that, within top economic journals, articles by men are cited less than those by women. More recently, Sherman and Tookes (2022) finds evidence consistent with these findings for papers published in finance. In all three of these cases, because women face obstacles to publishing in top economics and finance journals, female-authored papers that are eventually accepted for publication at those journals are ultimately of higher quality.

The estimates from the simplistic model thus seem suspicious. This result, however, may stem from the fact that citations are a noisy measure of quality and may themselves be biased based on gender of the inventor. It is this that our machine learning methodology next seeks to determine.

3.2 Adjusting for the Quality of Patents Using C-BERT

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.

With our two mappings, we were able to calculate the number of citations that would have been received if the lead inventor had been of the opposite gender. We used these calculations to generate the counterfactual number of citations for all patents, $\widehat{ForwardCitation}_i$, and plotted the updated histogram of citations in Panel B of Figure 4.[7].

[7]To address skewness and show more clearly, we the natural logarithm of citations and present a

When compared to Panel A, we can see that the C-BERT approach produces a histogram with a similar distribution.

To further consider the relative bias in citations at the patent-level, we need to compare the actual forward citation to the model-implied citations. Specifically, we calculate the following:

$$Bias_i = ForwardCitation_i - \widehat{ForwardCitation}_i. \tag{2}$$

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 difference would imply that a patent has received more citations than the quality-adjusted number suggested by opposite-gender model.

Plotting the difference between actual and model implied citations, Figure 5 visually suggests three important differences. First, and most importantly, we find that $Bias_i$ is negative for female lead inventors (plotted in red). The difference between female citations minus implied-by-male-model has a mean and median of -2.69 and -0.25, respectively. In contrast, for male lead inventor patents (plotted in blue), we see a difference with a mean and median of 1.10 and 0.13, respectively. Second, we find that these differences are skewed for both genders. Third, we find that the difference for any given patent is not strictly positive for males or negative for females. That is, as we would expect, some female lead inventors are over-cited, and some males are undercited.

We re-estimate the OLS in Equation 1, replacing the actual number of forward citations for a patent with $Bias_i$. As a reminder, a negative estimate of our coefficient of interest, β_1, suggests that women-led patents are undercited relative to the equivalent male-led patent.

The estimates in Table 3 present consistent evidence suggesting that patents with female lead inventors are undercited, relative to what would be expected had the patent

histogram in Panel B of Figure IA2

remained otherwise the same, save for the lead inventor instead being male. Interpreting our point estimate, relative to the sample mean of the number of forward citations, we find underciting of 27% for patents with female lead inventors. These results are 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.

Our estimates find consistent evidence suggesting that patents with female lead inventors would have been more likely to be highly cited if their lead inventors had been male. As reported in Table 4, the estimates suggest that representation in the top decile of cited patents by female lead inventors would have increased by 1.4—2.4 percentage points if the lead inventor had instead been male. These estimates are highly statistically significant and economically meaningful, suggesting an increase of 14% to 24%, relative to the sample mean, in the top decile (by construction, 10%).

These two results represent a first set of causal evidence suggesting that women 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 of Bias

4.1 Patent Category

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. To interpret the overall effects, we need to add the coefficients of the indicator for female lead inventor with the interaction term for each major category. Overall no clear pattern emerges from the data.

These estimates, however, mask differences in quality across male and female lead-inventor patents. 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. Using *Bias* as our dependent variable, column (2) of Table 5 shows that the unconditional forward citation data masks a significant disparity between female and male lead inventors, to the detriment of female lead-inventor patents. The summed coefficients

demonstrate a negative downward bias for female lead-inventor patents across all six major categories. 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 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 6 presents the interaction coefficients (Female Lead Inventor $\times$ Subcategory) graphically, with the estimates presented in Table IA4. 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. Column 1 of Table IA4 shows the estimates employing the raw citation counts, while column (2) presents the difference implied by the C-BERT model. The finer category classification exhibits somewhat more heterogeneity than the major classes.

Inspecting Figure 6, a number of clear patterns emerge from the estimates. First, once again, estimates from the model using raw citation counts as an outcome variable do not exhibit any clear differences between patents with male and female lead inventors. Once again, however, these estimates from the raw citation data mask important differences uncovered once we mediate for patent quality using the C-BERT model. For the vast majority of the technology subcategories, the estimates suggest that patents with lead female inventors are cited significantly less than they would be had they had a male lead inventor instead, and these citation undercounts are often substantial in magnitude. Overall the results are consistent with the baseline main results and suggest that (i)

patents with female lead inventors are overall of higher quality; (ii) these patents are undercited, relative to what they would have received had the lead inventor been male; and, (iii) as a result, raw counts of citations to male and female lead inventor patents look similar. These findings are also consistent with the selection effect predicted by the large literature on the obstacles to obtaining patents for female inventors. The estimates further underline the fact that using observed citations, without adjustments, may lead to misleading inferences about the quality of a patent.

4.2 Established Versus Emerging Fields

An interesting question is whether the patterns we see relate in some way to whether women are patenting in an established field versus in an emerging field of technology. To test this, we denote if the art unit first appeared within five years of the patent being granted. We then re-run our models, dropping the art-unit fixed effect and instead adding an indicator for an emerging field as well as an interaction between that indicator and the indicator for a female lead inventor.

The estimates are presented in Table 6. Our coefficient of interest is the interaction between the indicator for female inventors and emerging fields. Table 6 presents the models where the outcome variable is the difference in actual versus the predicted number of citations (for the opposite gender) from the C-BERT model. Estimates in this take suggest that new fields appear to receive more forward citations overall.

Mediating for the underlying quality of the patent a number of patterns emerge. First, our main effect is mirrored, with the indicator for female lead inventor loading significantly and suggesting that unconditionally patents with female lead inventors receive roughly 4.4 to 4.6 fewer citations than they would be expected if the lead inventor had been male. Second, patents in emerging fields appear to garner modestly more forward citations than those in established fields, depending on the specification, with approximately 0.06 to 0.9 more citations. Finally, and perhaps most interestingly, patents

in emerging fields with female lead inventors do not receive significantly fewer citations than would be expected if their lead inventor had been male. These estimates are insignificant for all specifications.

We hypothesize that that newer fields would not exhibit many barriers to entry for female inventors and researchers, given the lack of an established history of research and researchers. The results would be consistent with a lack of a selection effect, women do not face as large larger obstacles to patenting in newer fields, thus leading to a situation where patents with female inventors have the equal quality to those of male inventors. We leave further exploration of this finding to future research.

4.3 Corporate Innovations

An interesting aspect to consider is whether the pattern of undercitation we observed for female inventors also holds true for female authors of corporate innovations. Corporations generally have a strong interest in protecting their innovations and may therefore be more likely to ensure that their patents receive citations. To examine this possibility, we used data from Kogan, Papanikolaou, Seru, and Stoffman (2017) to identify patents that were associated with a particular firm, which we defined as corporate innovations.[8] We then interacted this indicator with a female lead indicator to investigate whether the undercitation of female inventors is more pronounced for corporate innovations.

The results shown in Table 7 offer several important insights. When we used Equation 2 as our dependent variable, we again found that female inventors received fewer citations than their male counterparts for equivalent patents. However, an interesting finding was that corporate innovations tended to receive more citations overall, with estimates ranging from 0.29 to 0.94 additional citations. However, when we looked specifically at corporate innovations with female lead inventors, we found that they received

[8] We use the expanded sample available from https://github.com/KPSS2017/Technological-Innovation-Resource-Allocation-and-Growth-Extended-Data.

fewer citations than their male counterparts. These results suggest that the undercitation of female inventors may be even more pronounced for corporate innovations.

4.4 Time After the Patent is Granted

To gain a deeper understanding of the pattern of undercitation for female inventors, we also considered the timing of the bias in citations. One possibility is that the bias is present from the outset but diminishes over time as people become more familiar with the patent. 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 the bias in citations, we created separate samples based on the number of years that had passed since a patent was granted. Specifically, we divided the sample into four periods: [0-1) years, [1-5) years, [5-10) years, and [10-20] years. Using these subsamples, we re-ran our C-BERT methodology to estimate the relative bias in citations after controlling for patent quality. This allowed us to examine whether the bias in citations is present from the outset or if it changes over time.

Upon analyzing the citation patterns based on the number of years that have passed since the patent was granted, we found that the bias against female inventors tends to increase over time. The estimates in Table 8 show that the bias is statistically insignificant in the first period ([0-1) years), but becomes more pronounced 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. These findings suggest that the undercitation of female inventors may be self-reinforcing and become more difficult to overcome as time goes on.

Our estimates, which are shown in Table 8, are consistent with our baseline specification. When we sum the point estimates across the four periods, we obtain results

that are similar 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.

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 9 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 -0.002 citations to -0.017 citations, with only one 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 -0.012 citations to -0.051, with again only one 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 10. 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 1.2 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 do not argue 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 Patent Gender

One potential explanation for our findings is that the way in which we assigned gender to patents may have influenced the results. In our baseline method, we used the name of the first inventor to determine the gender of patents with multiple inventors. This is a common practice in the literature, but it is worth noting that it is possible that examiners and inventors may consider all inventors and not just the first author when assigning

gender. Therefore, it is possible that our results are partially driven by this method of gender assignment.

To address the potential issue of our method of gender assignment influencing our results, we tried two different approaches. The first approach involved limiting our sample to patents with only one author and re-running our C-BERT model. This allows us to eliminate any concerns that the gender of additional authors, other than the lead author, may be driving the baseline results.

Although we altered our sample selection by including only patents with a single author, we found that the results were similar to those obtained in the baseline approach. When we re-ran our baseline specification, Equation 1, using the single-author sample, we still found a statistically significant bias against female inventors. The estimates shown in Panel B of Table IA3 were consistent with those in Table 3. These findings suggest that the gender of the lead inventor alone is sufficient to capture the gender of the patent, at least in part.

In order to examine the role of gender while controlling for the potential impact of single-author patents, we constructed a new sample that included patents with inventors of the same gender, both single authors and teams.[9] When we re-ran our C-BERT model using this expanded sample, we found that female inventors were still undercited. The estimates shown in Panel C of Table IA3 indicate that the bias increased from 2.2 to 3.7, which suggests that the undercitation of female inventors may be even more pronounced when considering patents with multiple authors.

6.2 Patents with Zero Citations

Our analysis to this point has focused on patents that received at least one citation. While this allows us to examine the extent of discrimination in the number of citations

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

received, it is possible that female inventors may not only receive fewer citations, but also be completely overlooked and receive zero citations. If this is the case, our analysis may underestimate the true level of bias that female inventors face, since it only considers patents that received a positive number of citations. To fully understand the extent of discrimination against female inventors, it may be necessary to also consider patents that received no citations, regardless of their quality.

To address the potential issue of missing information on patents that received no citations, we modified our sample to include all patents, including those that did not receive any citations. Using this expanded sample, we re-ran our C-BERT methodology. This allowed us to incorporate information on patents that received no citations in our counterfactual analysis, providing a more comprehensive understanding of discrimination against female inventors.

Re-estimating our baseline specification using the expanded sample that includes patents with zero citations, we find evidence that female inventors also face biases in their chances of receiving any citations at all. The estimates in Panel A of Table IA3 indicate that the bias against female inventors is even more pronounced when considering patents that received no citations. These findings suggest that female inventors may face obstacles in obtaining citations not just in comparison to male inventors, but also in absolute terms.

6.3 Overfitting 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 IA3, 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

It is worth considering the relationship between the economic importance of a patent, as evaluated by public markets, and the bias in patent citations. As we have previously discussed, the biases in citations tend 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. This suggests that the economic value of a patent may not be closely tied to the biases in citations that are observed over time.

To investigate the relationship between private economic value and citations, we used the patent-level measure of private economic value proposed in Kogan et al. (2017). This measure is based on data on patent issues for U.S. firms, which is combined with the stock market's response to news about patents. By using this measure, we were able to examine the relationship between private economic value and citations at the patent level. Specifically, we use the log value of innovation, deflated to 1982 (million) dollars using the CPI.

The results shown in Panel A of Table 11 suggest that there is a negative relation-

ship between the economic value of a patent and the bias in citations. Specifically, the estimates in this table, which are based on the bias in citations defined by Equation 2, indicate that patents that are undercited tend to have a higher economic value, while patents that are overcited tend to have a lower economic value. This suggests that markets may be able to rationally attribute value to patents and take the biases in citations into account when evaluating their economic importance.

To examine the relative importance of our C-BERT measure of citation bias versus forward citations for patents, we decomposed our bias measure into both $ForwardCitation_i$ and $\widehat{ForwardCitation_i}$ and compared the two. The results shown in Panel B of Table 11 indicate that the relationship between $\widehat{ForwardCitation_i}$ and the economic value of a patent is stronger. This finding suggests that observed forward citations may be noisy and may need to be adjusted to account for possible biases in citing patterns. It also highlights the value of our C-BERT measure as a tool for identifying and correcting such biases.

8 Discussion and Conclusion

We provide causal evidence that patents with female lead inventors are undercited, relative to what they would have happened had their patent had a male lead inventor. Importantly, this effect is masked in raw citation count data, where true patent quality cannot be mediated. 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.

There are important economic implications to our findings. First, if female inventors are undercited, relative to male peers with equivalent patents, and their compensation for their innovative labor is accordingly harmed, this may discourage women from en-

tering the innovation economy.[10] 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 prior findings may be warranted.

Our paper also makes an important methodological contribution to the finance and economic literatures by introducing the C-BERT methodology for causal inference. Finance 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.

[10]The literature has highlighted that innovative activity is motivated by expected profits derived from the property rights granted to the patentee, Moser (2005, 2013).In related research, the marginal investor values patents Aghion et al. (2013); Hall et al. (2005); Hirschey and Richardson (2004); Hirshleifer et al. (2013).

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, 2012, Why don't women patent?, Technical report, National Bureau of Economic Research.

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.

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.

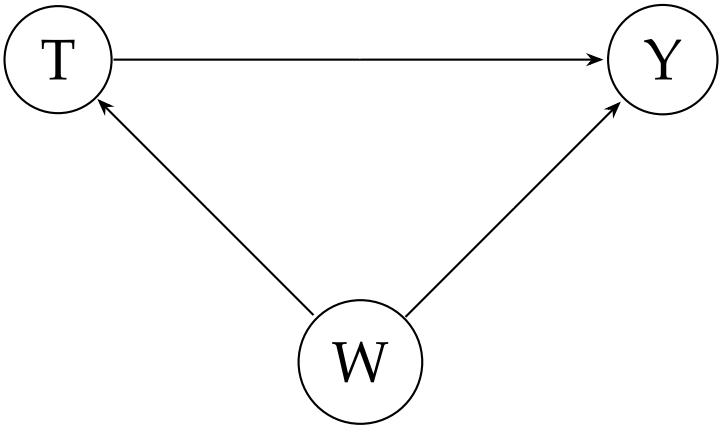
Sherman, Mila Getmansky, and Heather E Tookes, 2022, Female representation in the academic finance profession, *Journal of Finance* 77, 317–365.

Subramani, Gauri, Abhay Aneja, and Oren Reshef, 2021, Persistence and the gender innovation gap, Working paper.

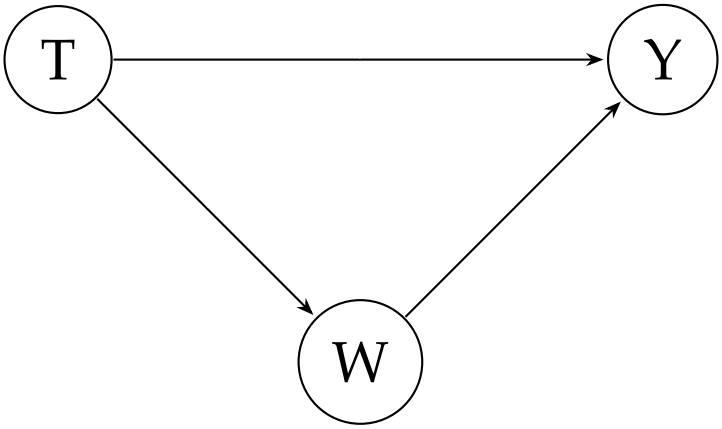
Tzioumis, Konstantinos, 2018, Demographic aspects of first names, *Scientific Data* 5, 1–9.

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.



PANEL A: MODEL FOR ATT, CONFOUNDING



PANEL B: MODEL FOR NDE, MEDIATING

FIGURE 1: C-BERT

This figure replicates Figure 1 of Veitch et al. (2020) and illustrates they assumption of our mediating approach. Y is the outcome of interest, T is the treatment, and W are the sequence of words. Panel A depicts the average treatment effect (ATT), with the assumption that W 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.

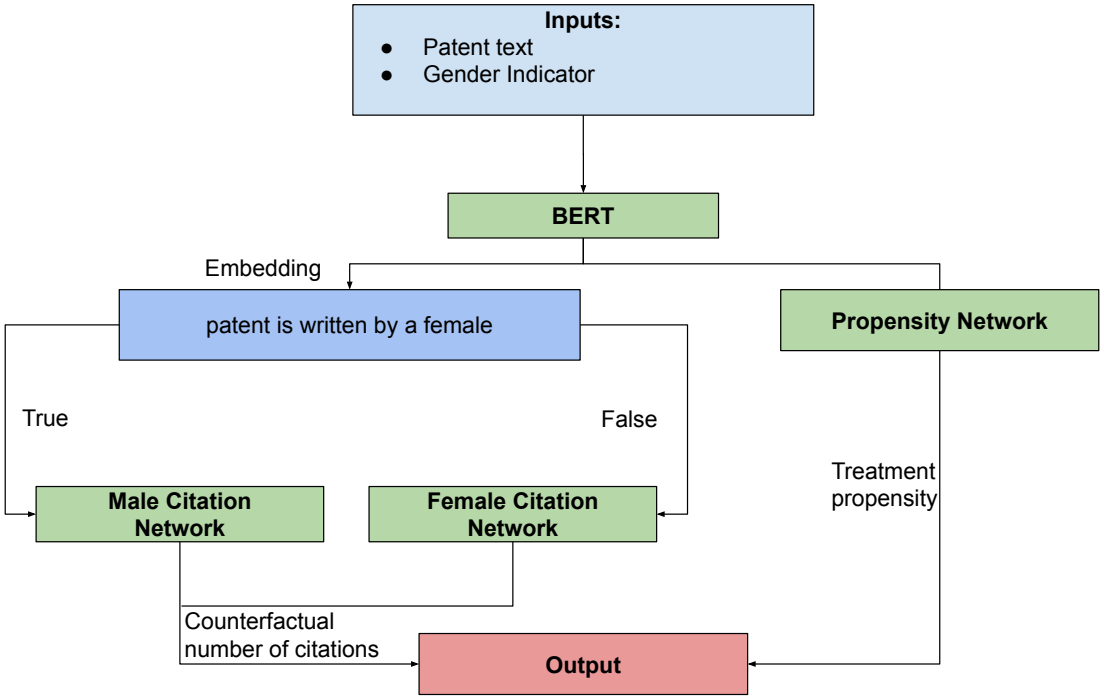


FIGURE 2: C-BERT ESTIMATION PROCEDURE

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.

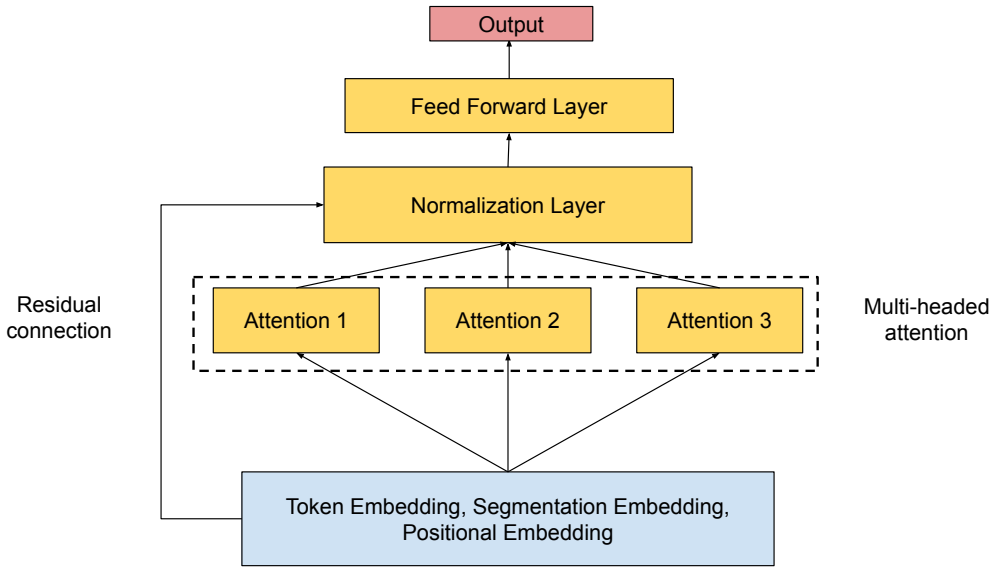
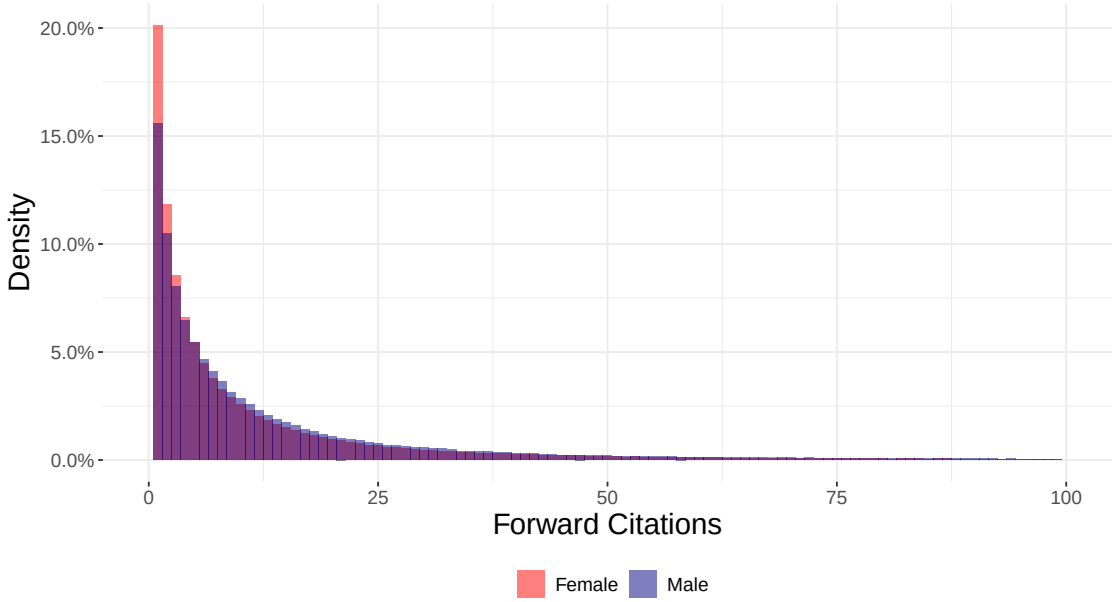
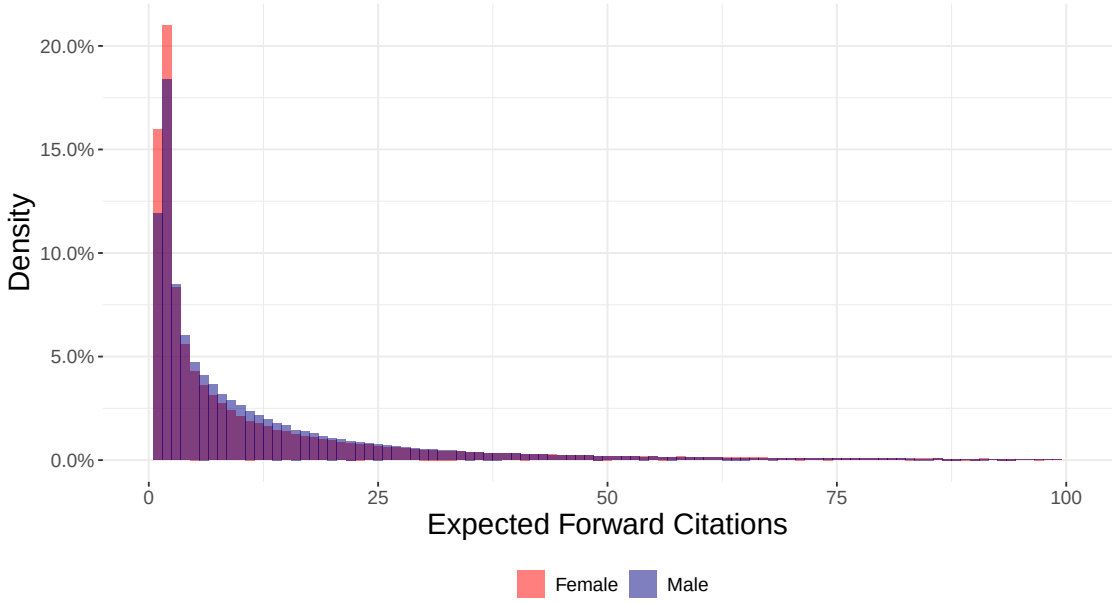


FIGURE 3: ENCODER MODULE

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.



(a) Panel A: Observed Forward Citations



(b) Panel B: Model Implied Forward

FIGURE 4: 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 IA2.

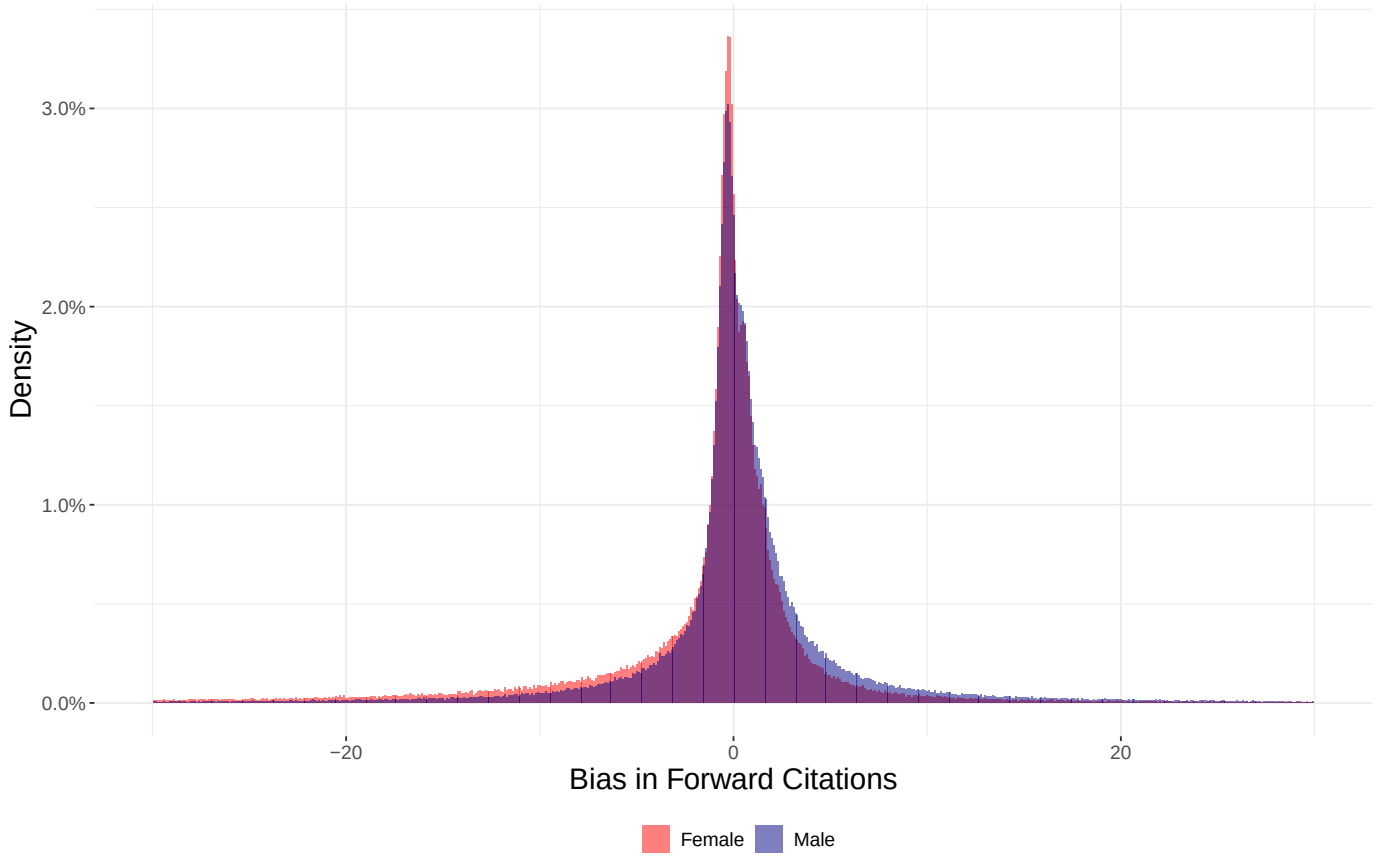


FIGURE 5: UNDERCITATION OF FEMALE LEAD PATENTS

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.

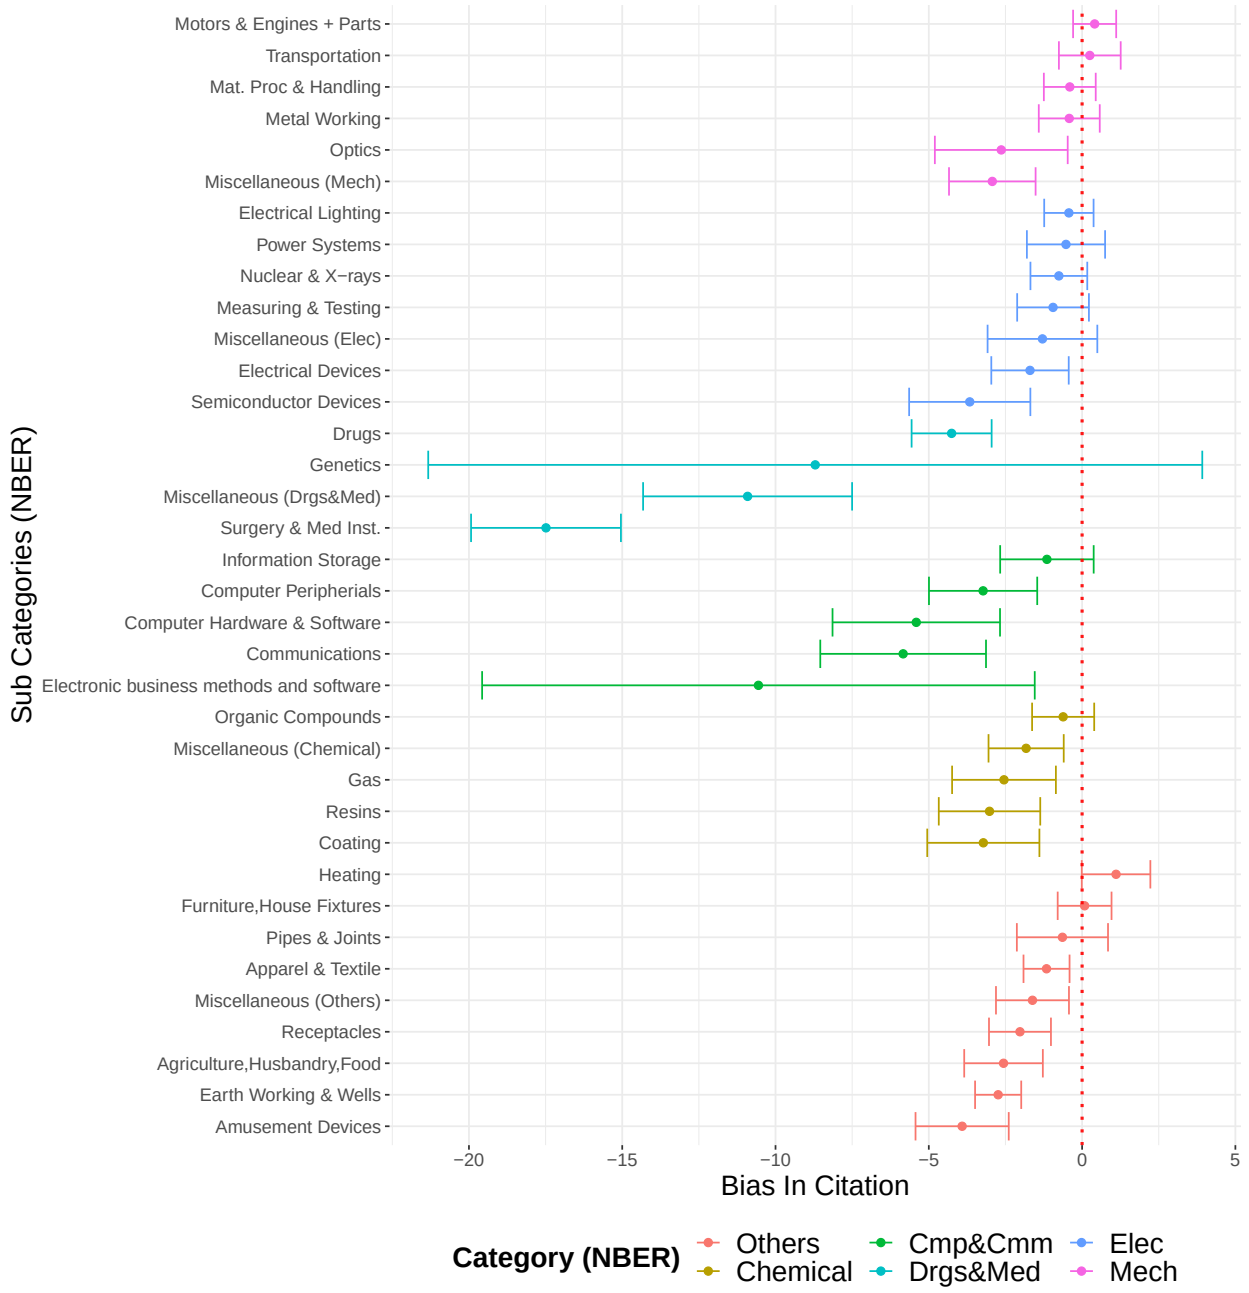


FIGURE 6: BIAS IN CITATIONS BY PATENT SUB CATEGORIES

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.

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, BY GENDER

This table reports estimates of Equation 1 and studies the number of forward citations by the gender of the lead inventor. The estimate uses the observed level of forward citations as its dependent variable. 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	-2.918^{***}	-2.113^{***}	-1.532^{***}	-1.167^{***}
	(0.668)	(0.358)	(0.189)	(0.123)
Intercept	20.939^{***}			
	(1.581)			
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	408,707	408,707	408,707	408,707
Adjusted R^2	0.001	0.086	0.171	0.270

TABLE 3: ACTUAL MINUS EXPECTED FORWARD CITATIONS

This table estimates Equation 1 and studies the difference in the number of citations by the gender of the lead inventor. 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.

	Bias in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-4.600^{***}	-4.649^{***}	-4.611^{***}	-4.651^{***}
	(0.428)	(0.418)	(0.385)	(0.229)
Intercept	1.316^{***}			
	(0.125)			
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	408,707	408,707	408,707	408,707
Adjusted R^2	0.007	0.012	0.057	0.252

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: Observed Patent Citations

	Top Decile Patent			
	(1)	(2)	(3)	(4)
Lead Inventor Female	-0.018^{***}	-0.013^{***}	-0.011^{***}	-0.008^{***}
	(0.005)	(0.002)	(0.001)	(0.001)
Intercept	0.127^{***}			
	(0.012)			
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	408,707	408,707	408,707	408,707
Adjusted R^2	0.001	0.094	0.128	0.051

Panel B: Model Adjustment

	Model Expected Forward Citation in Top Decile			
	(1)	(2)	(3)	(4)
Lead Inventor Female	0.014^{*}	0.019^{***}	0.021^{***}	0.024^{***}
	(0.007)	(0.005)	(0.004)	(0.001)
Intercept	0.116^{***}			
	(0.011)			
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	408,707	408,707	408,707	408,707
Adjusted R^2	0.0005	0.092	0.126	-0.062

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	Bias in Forward Citations
	(1)	(2)
Female Lead Inventor	−0.151	−3.430***
	(0.358)	(0.107)
Chemical × Female Lead Inventor	−0.721**	−0.488**
	(0.306)	(0.221)
Computers and Communication × Female Lead Inventor	−1.271***	−3.150***
	(0.261)	(0.921)
Drugs and Medical × Female Lead Inventor	−5.817***	−7.517***
	(0.392)	(0.444)
Electrical × Female Lead Inventor	−0.765**	0.048
	(0.368)	(0.211)
Mechanical × Female Lead Inventor	−0.031	0.659***
	(0.265)	(0.141)
NBER Category FE	Yes	Yes
Examiner x Customer FE	Yes	Yes
Examiner Art Unit FE	Yes	Yes
Patent Issue Year FE	Yes	Yes
Observations	408,706	408,706
Adjusted R^2	0.271	0.253

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.

	Bias in Forward Citations			
	(1)	(2)	(3)	(4)
New Field	0.058	0.147	0.400	0.896**
	(0.172)	(0.419)	(0.355)	(0.368)
Lead Female Inventor	−4.429***	−4.502***	−4.493***	−4.587***
	(0.486)	(0.473)	(0.438)	(0.231)
New Field × Lead Female Inventor	−0.580	−0.513	−0.421	−0.335
	(0.539)	(0.525)	(0.495)	(0.225)
Intercept	1.301***			
	(0.136)			
Patent Issue Year FE	No	Yes	Yes	Yes
Customer FE	No	No	Yes	No
Examiner FE	No	No	Yes	No
Examiner x Customer FE	No	No	No	Yes
Observations	408,707	408,707	408,707	408,707
Adjusted R^2	0.007	0.012	0.057	0.252

TABLE 7: DIFFERENCE IN FORWARD CITATIONS, CORPORATE INNOVATION

This table studies the citation patterns of corporate innovations. *Corporate Innovation* takes the value of one if the patent is associated with a company in CRSP. The dependent variable is the difference in 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.

	Bias in Forward Citations			
	(1)	(2)	(3)	(4)
Corporate Innovation	0.290	0.609*	0.605**	0.938***
	(0.334)	(0.325)	(0.297)	(0.230)
Lead Female Inventor	−4.048***	−4.117***	−4.076***	−3.823***
	(0.389)	(0.389)	(0.340)	(0.150)
Corporate Innovation × Lead Female Inventor	−1.368**	−1.336**	−1.336**	−2.165***
	(0.583)	(0.588)	(0.570)	(0.434)
Intercept	1.204***			
	(0.227)			
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	408,707	408,707	408,707	408,707
Adjusted R^2	0.007	0.012	0.057	0.252

TABLE 8: 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, 5-10, and 10-20 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.

	Bias in Forward Citations			
	0-1 Years	2-5 Years	5-10 Years	10-20 Years
	(1)	(2)	(3)	(4)
Lead Inventor Female	0.048	-0.706^{***}	-1.294^{***}	-2.683^{***}
	(0.029)	(0.023)	(0.072)	(0.138)
Examiner x Customer FE	Yes	Yes	Yes	Yes
Art Unit FE	Yes	Yes	Yes	Yes
Patent Issue Year FE	Yes	Yes	Yes	Yes
Observations	8,449	185,654	236,652	210,698
Adjusted R^2	0.622	0.370	0.455	0.446

TABLE 9: 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

	Bias in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-0.017^{**}	-0.012	-0.002	-0.009
	(0.008)	(0.009)	(0.008)	(0.010)
Intercept	0.162^{***}			
	(0.026)			
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	57,556	57,556	57,556	57,556
Adjusted R^2	0.00001	0.013	0.042	-0.008

Panel B: Citation Added by Male Lead Examiners

	Bias in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-0.018	-0.012	-0.022	-0.051^{***}
	(0.035)	(0.037)	(0.032)	(0.014)
Intercept	0.118^{***}			
	(0.040)			
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	146,101	146,101	146,101	146,101
Adjusted R^2	0.00000	0.008	0.027	0.085

TABLE 10: 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

	Bias in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-0.483^{***}	-0.463^{***}	-0.481^{***}	-0.483^{***}
	(0.061)	(0.063)	(0.042)	(0.065)
Intercept	0.667^{***}			
	(0.034)			
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	25,016	25,016	25,016	25,016
Adjusted R^2	0.003	0.014	-0.001	-0.029

Panel B: Citation Added by Male Lead Inventors

	Bias in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	-1.528^{***}	-1.530^{***}	-1.485^{***}	-1.207^{***}
	(0.164)	(0.155)	(0.141)	(0.089)
Intercept	1.264^{***}			
	(0.137)			
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	231,082	231,082	231,082	231,082
Adjusted R^2	0.002	0.010	0.093	0.139

TABLE 11: 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). Panel A uses the difference in forward citations, as defined in Equation 2, as its main independent variable. Panel B uses the observed forward citation and expected forward citations as its independent variables. 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: Difference in Forward Citations

	$\log(dollar)$			
	(1)	(2)	(3)	(4)
Bias in Forward Citations	−0.0004	−0.0004**	−0.0005***	−0.001**
	(0.0003)	(0.0002)	(0.0001)	(0.0002)
Constant	0.682***			
	(0.135)			
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	159,138	159,138	159,138	159,138
Adjusted R^2	0.00002	0.120	0.510	0.381

Panel B: Decomposition of Citation Measures

	$\log(dollar)$			
	(1)	(2)	(3)	(4)
Forward Citation	−0.0001	−0.0002	0.0001	−0.00003
	(0.0003)	(0.0003)	(0.0002)	(0.0002)
$\widehat{\text{Forward Citation}}$	0.005***	0.004***	0.002***	0.003***
	(0.0004)	(0.0003)	(0.001)	(0.001)
Female Lead Inventor	0.097***	0.018	0.015*	0.003
	(0.029)	(0.024)	(0.008)	(0.006)
Intercept	0.607***			
	(0.142)			
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	161,737	161,737	161,737	161,737
Adjusted R^2	0.011	0.140	0.519	0.398

INTERNET APPENDIX
FOR ONLINE PUBLICATION

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 IA1, 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.

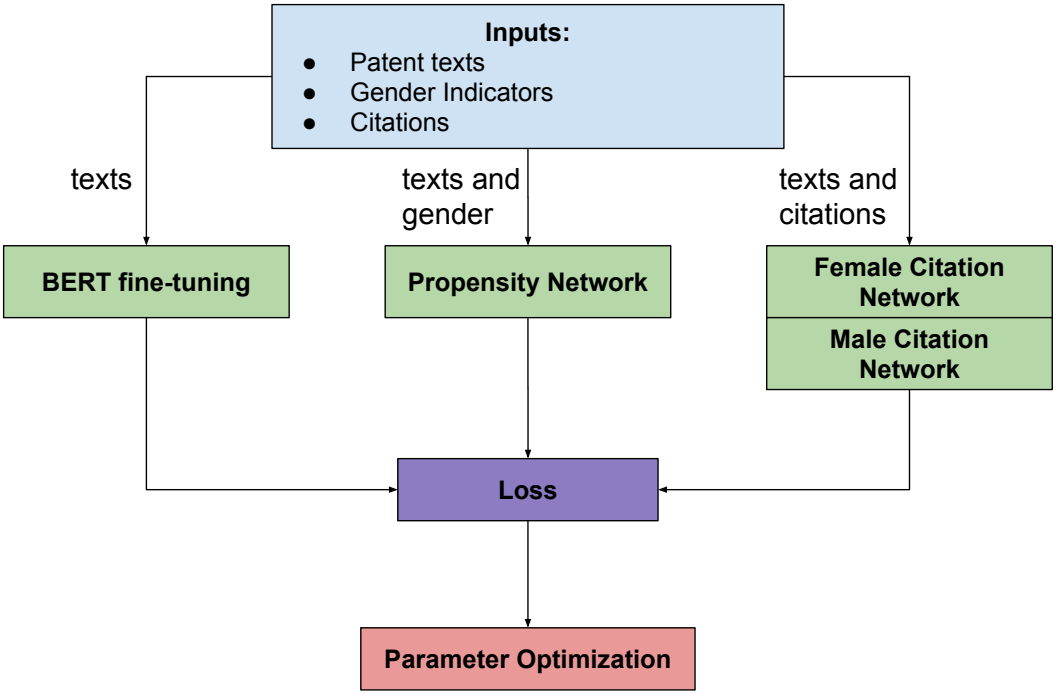
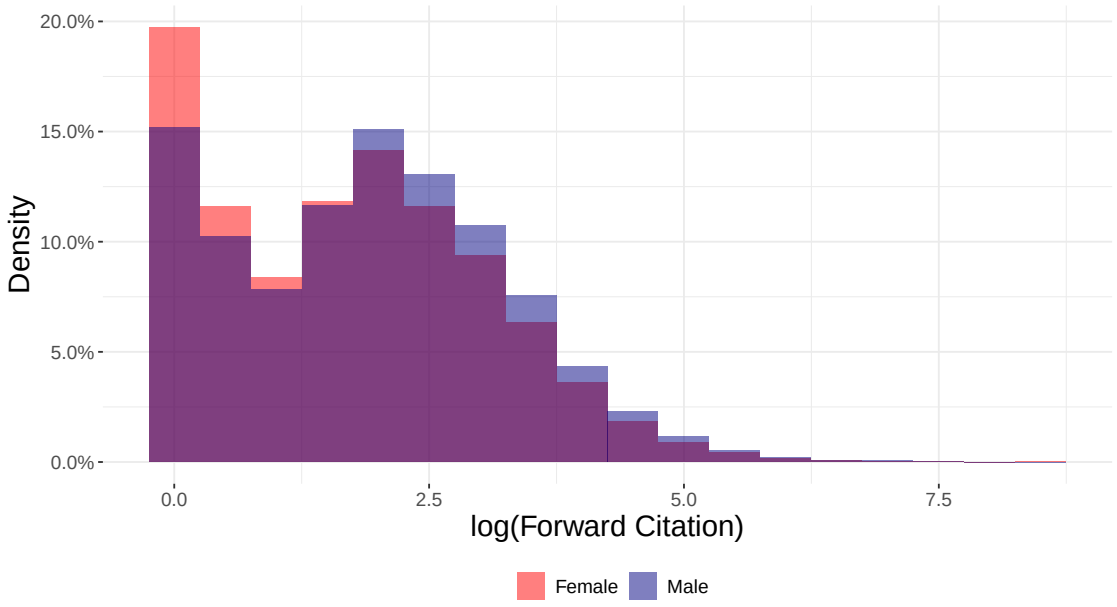
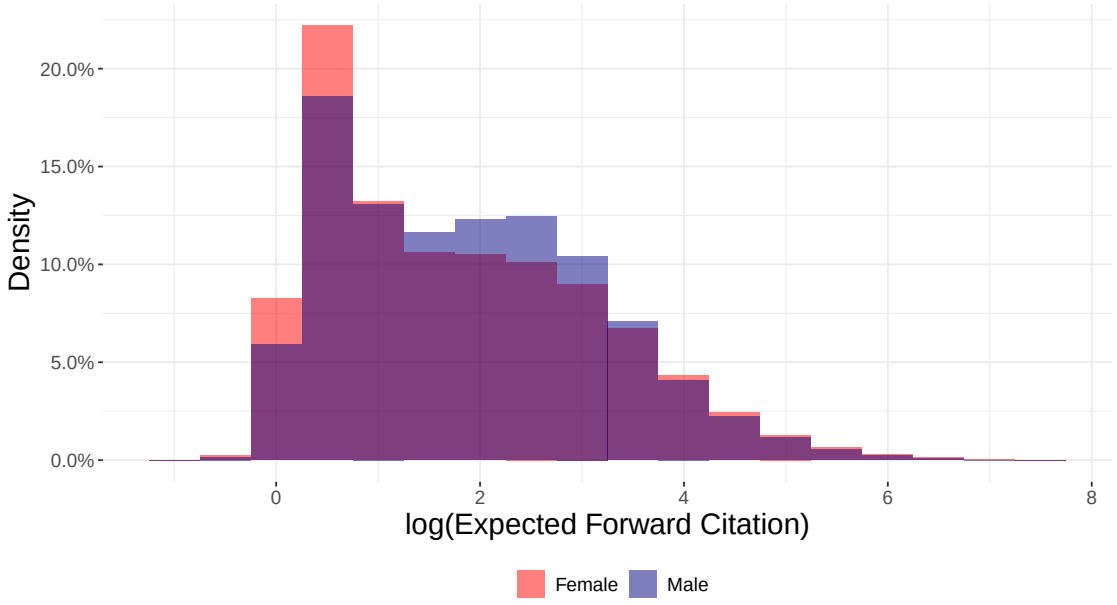


FIGURE IA1: C-BERT TRAINING PROCEDURE

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.



(a) Forward Citations



(b) Model Implied Forward

FIGURE IA2: DISTRIBUTION OF FORWARD CITATIONS

This figure illustrates the transformation from forward citations to expected forward citations. Panel A uses the natural logarithm of forward citations while Panel B uses the natural logarithm of forward citations expected from our model. The vertical axis in both panels measures the percent of the distribution. Red bars correspond to females, blue bars correspond to males, and purple bars correspond to the overlapping region.

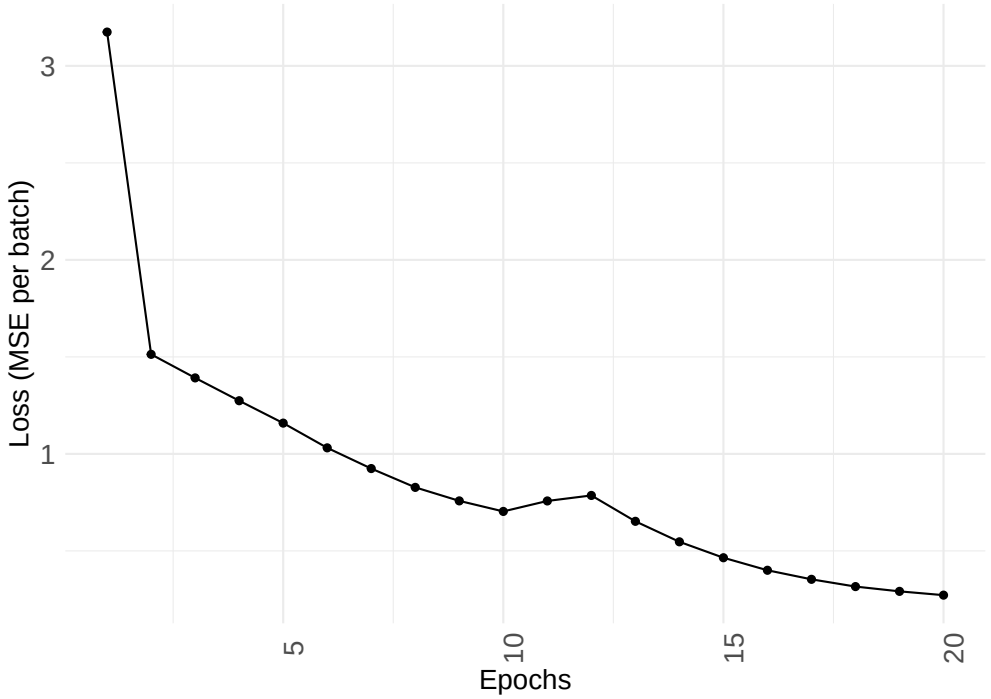


FIGURE IA3: LOSS FUNCTION

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: 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	Difference in Forward Citations
	(1)	(2)
Lead Female Inventor	-1.833^{***}	-4.324^{***}
	(0.183)	(0.192)
Electricity × Lead Female Inventor	0.485	-0.393^{*}
	(0.288)	(0.205)
Fixed Constructions × Lead Female Inventor	0.744^{**}	1.771^{***}
	(0.366)	(0.181)
Human Necessities × Lead Female Inventor	-1.446^{**}	-4.736^{***}
	(0.539)	(0.199)
Mechanical Engineering × Lead Female Inventor	2.084^{***}	3.143^{***}
	(0.209)	(0.290)
Performing Operations × Lead Female Inventor	1.226^{***}	0.979^{***}
	(0.096)	(0.138)
Physics × Lead Female Inventor	0.889^{***}	-0.494
	(0.200)	(0.445)
Textiles × Lead Female Inventor	3.376^{***}	1.271^{**}
	(0.542)	(0.492)
CPC Section FE	Yes	Yes
Examiner x Customer FE	Yes	Yes
Examiner Art Unit FE	Yes	Yes
Patent Issue Year FE	Yes	Yes
Observations	408,679	408,679
Adjusted R^2	0.272	0.253

TABLE IA3: ROBUSTNESS TO SAMPLE SELECTION

This table establishes robustness of our baseline specification in Table 3. Panel A expands our sample to include patents that have never been cited. Panel B uses a single author patent. Panel C uses both single-author patents and patents where all inventors share the same gender. The sample covers patents issued from 1976-01-01 through 2008-12-31. 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: All Patents

	Difference in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	−6.018***	−6.146***	−6.152***	−6.698***
	(0.804)	(0.742)	(0.714)	(0.386)
Intercept	0.342***			
	(0.068)			
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	574,621	574,621	574,621	574,621
Adjusted R^2	0.011	0.019	0.042	0.130

Panel B: Single Author

	Difference in Forward Citations			
	(1)	(2)	(3)	(4)
Lead Female Inventor	−2.526***	−2.576***	−2.438***	−2.229***
	(0.283)	(0.268)	(0.205)	(0.105)
Intercept	0.113			
	(0.241)			
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	100,130	100,130	100,130	100,130
Adjusted R^2	0.002	0.014	0.116	0.343

Panel C: Inventors Same Gender

	Difference 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	102,864	102,864	102,864	102,864
Adjusted R^2	0.003	0.016	0.192	0.378

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 *Bias* 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. The sample covers patents issued from 1985-01-01 through 1995-12-31. 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	Bias in Forward Citations
	(1)	(2)
Female Lead Inventor	0.606	−1.877***
	(0.464)	(0.465)
Agriculture,Husbandry,Food× Female Lead Inventor	0.051	−2.563***
	(0.624)	(0.654)
Amusement Devices× Female Lead Inventor	1.467	−3.913***
	(2.895)	(0.776)
Apparel & Textile× Female Lead Inventor	2.463***	−1.160***
	(0.660)	(0.383)
Coating× Female Lead Inventor	−1.570	−3.221***
	(1.137)	(0.933)
Communications× Female Lead Inventor	−2.293***	−5.837***
	(0.596)	(1.379)
Computer Hardware & Software× Female Lead Inventor	−1.953**	−5.409***
	(0.893)	(1.394)
Computer Peripherials× Female Lead Inventor	0.343	−3.229***
	(0.973)	(0.901)
Drugs× Female Lead Inventor	−5.078***	−4.256***
	(0.734)	(0.666)
Earth Working & Wells× Female Lead Inventor	−2.054***	−2.738***
	(0.723)	(0.384)
Electrical Devices× Female Lead Inventor	−1.339***	−1.699**
	(0.493)	(0.644)
Electrical Lighting× Female Lead Inventor	−2.472**	−0.431
	(1.034)	(0.411)
Electronic business methods and software× Female Lead Inventor	−10.009***	−10.557**
	(3.557)	(4.598)
Furniture, House Fixtures× Female Lead Inventor	−0.564	0.082
	(0.561)	(0.448)
Gas× Female Lead Inventor	−0.750	−2.548***
	(0.849)	(0.864)
Genetics× Female Lead Inventor	−9.398*	−8.704
	(5.058)	(6.441)
Heating× Female Lead Inventor	−1.962***	1.108*
	(0.716)	(0.571)
Information Storage× Female Lead Inventor	−2.417***	−1.147
	(0.873)	(0.778)

continued on next page...

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

	Dependent variable:	
	Forward Citations	Bias in Forward Citations
	(1)	(2)
Mat. Proc & Handling× Female Lead Inventor	0.247	−0.402
	(0.699)	(0.432)
Measuring & Testing× Female Lead Inventor	−0.764	−0.947
	(0.541)	(0.597)
Metal Working× Female Lead Inventor	0.301	−0.419
	(0.997)	(0.507)
Miscellaneous (Chemical)× Female Lead Inventor	−1.369***	−1.826***
	(0.435)	(0.626)
Miscellaneous (Drgs&Med)× Female Lead Inventor	−6.502**	−10.911***
	(2.982)	(1.739)
Miscellaneous (Elec)× Female Lead Inventor	−0.740	−1.293
	(0.627)	(0.913)
Miscellaneous (Mech)× Female Lead Inventor	−2.460*	−2.929***
	(1.252)	(0.720)
Miscellaneous (Others)× Female Lead Inventor	−0.730	−1.619**
	(0.904)	(0.607)
Motors & Engines + Parts× Female Lead Inventor	−0.112	0.409
	(0.694)	(0.359)
Nuclear & X-rays× Female Lead Inventor	−3.266***	−0.758
	(0.870)	(0.473)
Optics× Female Lead Inventor	−2.843*	−2.636**
	(1.640)	(1.105)
Organic Compounds× Female Lead Inventor	−2.419**	−0.618
	(1.092)	(0.517)
Pipes & Joints× Female Lead Inventor	−5.196***	−0.641
	(0.920)	(0.759)
Power Systems× Female Lead Inventor	−1.033*	−0.525
	(0.557)	(0.651)
Receptacles× Female Lead Inventor	−2.757***	−2.026***
	(0.781)	(0.515)
Resins× Female Lead Inventor	−1.996**	−3.020***
	(0.835)	(0.844)
Semiconductor Devices× Female Lead Inventor	−2.342**	−3.665***
	(0.878)	(1.009)
Surgery & Med Inst.× Female Lead Inventor	−7.304***	−17.487***
	(1.349)	(1.247)
Transportation× Female Lead Inventor	−0.859	0.251
	(0.541)	(0.514)
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	408,706	408,706
Adjusted R^2	0.275	0.255