Understanding Causal Patenting Disparities: An AI + Human Perceptions Approach*

Lin William Cong† Stephen Yang‡

November 2024

Abstract

Discrimination in innovation and entrepreneurship harms both economic efficiency and social mobility. In patenting, a large gap remains to be bridged both in terms of patenting policies and structural disparities. Building on causal machine learning and ppplying large-language-model embeddings to create objective measures of patent quality, we document systematic bias against under-represented groups by the United States Patent & Trademark Office, unexplained by gaps in the purported criteria for patent approvals. Manually labeled human perception tests and decreased gendered discrimination in API applicants corroborate that inventors' names illegally influence examiners' decisions. In jointly investigating multiple dimensions of discrimination for the first time, we discover a novel *affiliation bias*, with large- and private-firm-affiliated inventors gaining unjustified advantages significantly greater than those from gendered/racial discrimination and not explained by homophily. Further, various factors from affiliation characteristics to innovation quality to location can either mitigate or compound disparities, with misogyny persisting even with advantageous affiliations. We rationalize the findings in a unifying theoretical framework with correlation neglect driving examiner discrimination in which examiners mistakenly interpret traits such as gender, race, and affiliation as providing additional information.

JEL Classification: G30, J15, J16, O31.

Keywords: Causal ML, Gender/Racial Disparities, Innovation, LLMs.

*We thank Josh Lerner for helpful comments and discussions. We are grateful to Tanvi Kiggar, Rishi Kumar, and Jordan Velte for outstanding research assistance.

†Cornell SC Johnson College of Business and NBER. Email: will.cong@cornell.edu

‡Pace Academy. Email: stephen.yang25@paceacademy.org.

1. Introduction

Racial and gendered disparities in economic outcomes have been widely documented and acknowledged as one of the most pressing concerns in the world (e.g., Bertrand, Kamenica, and Pan, 2015; Blau and Kahn, 2017; Chetty et al., 2014, 2020; Derenoncourt and Montialoux, 2021). These differences not only are inherently unfair and unjust but also hinder economic growth and productivity, with a particularly negative impact on talent allocation (Altonji and Blank, 1999; Bertrand, 2011; Hsieh et al., 2019). As the embodiment of the oft-lauded "American Dream," innovation and entrepreneurship ought to act as a bridge between social groups.

However, this American Dream seems to be, at least historically, a false hope for most groups, due to the extensive and concerning disparities in VC investment opportunities as well as loan availability (Chernenko and Scharfstein, 2022; Fairlie, Robb, and Robinson, 2022; Fairlie and Robinson, 2023; Wang, Wu, and Hitt, 2023). This paper seeks to fill a void in this literature: the documentation and identification of unfair disparities in the patenting process. In the past, studies have investigated gaps in patenting activity between races and genders.[1] However, while some papers such as Hochberg et al. (2023) perform causal analyses on outcomes, the literature contains no unifying comparative analysis of the *relative* causes and effects of different forms of patenting discrimination.

This is an essential question to answer, given that innovation has the potential to either exaggerate or ameliorate social disparities. Patenting is a critical component of innovation, especially apropos social mobility (Aghion et al., 2019; Akcigit, Grigsby, and Nicholas, 2017). Determining a suitable policy response to the undisputable existence of disparities requires acknowledging that there are two potential channels for differences in patenting rates: first is lower access to education and funding leading to, in aggregate, worse innovation produced by certain groups. Second is a "discrimination" effect, in

[1] See below for an in-depth discussion of the existing literature.

which minority inventors may receive worse outcomes *ceteris paribus*, i.e., similar to discrimination in job or loan applications. While both situations are equally unjust, they demand substantially different policies from the perspective of the social planner and the government. We seek to isolate and quantify the relative effects of these two mechanisms.

First, we provide novel, comprehensive statistics on acceptance rate and acceptance quality across races and genders. In general, white and Asian & Pacific Islanders (henceforth, API) inventors receive the highest acceptance rates, while black and Hispanic inventors receive lower acceptance rates. In terms of application quality, patents filed by API inventors exhibit the highest average quality, whereas white patents have somewhat higher quality than black and Hispanic inventors. These differences are significantly lower in magnitude than those in acceptance rates, indicating that there are potentially two contributing factors to the under-acceptance of minority patents: first, perhaps due to an educational gap or socioeconomic disparities, they produce slightly lower quality innovation on average. However, this does not explain the majority of the difference, which instead seems to be a result of discriminatory treatment. Given the potential contemporary impact of such a gap, it is thus of utmost importance to identify the extent to which minorities are treated unfairly.

In the case of disparities in job opportunities and hiring, the seminal Bertrand and Mullainathan (2004) conducts a field experiment to test for racial discrimination. By submitting identical resumes with different, stereotypically black or white, names in job applications, they find significant differences in the number of callbacks and interview offers given to the "white" and "black" candidates. There are likely two reasons why such a study does not exist in the case of patents. First, a similar field test that provides clear evidence of disparate treatment would be infeasible.[2] Second, the area of patent applications, which are the ex-ante component of patent grants and thus critical to the patent

[2]Either submitting large quantities of fictional patent applications to the real USPTO or hiring a significant number of real patent examiners to evaluate applications would be prohibitively expensive and, in the case of the former, likely illegal.

acceptance process, has been systematically understudied, partially because of prior data constraints. We present solutions to both of these problems by leveraging a Generative AI measure of patent application quality similar to that introduced in Yang (2024). Applying Large-Language Model technology and machine learning to the patent's contents, we train an objective measure of application quality on historical data to use as a control. Fundamentally, we then regress patent acceptances on demographic characteristics while controlling for the quality and content of the application, thus determining the unique effect of demographics. This quality measure is a powerful predictor for other patent outcomes and ought to capture the sole determinants of patent acceptance given the USPTO criteria of originality, non-triviality, and usefulness—the persistence of demographics as predictors indicates bias.

Through this process, we determine a plausibly causal relation between race and patent application acceptances. Even when controlling for application quality, applications filed by Black and Hispanic inventors are much less likely to be accepted than those submitted by white and API inventors, with Black inventors experiencing the worst discrimination at 3 percentage points relative to whites. In order to determine the mechanism behind these disparities, we study cross-sectional variation across types of applicants. Based on the assignee (legal owner of the patent's intellectual property rights), we label the patents as filed by either individual inventors, private firms or public firms. If a subtle mechanism such as the credibility of an inventor in the eyes of the examiner (as opposed to intentional discrimination) is at play, certain other factors may ameliorate the effect. For example, if a subconsciously racist examiner reviews a patent filed by a minority inventor working for Ford or Pfizer, they may be more prone to overlook their racial bias in favor of bias toward the company name. If examiners were intentionally racist, the effect would instead remain the same across all subsets.

We find that there are three levels of discrimination across the different samples. Indeed,

in the case of public firms, there is much less of a disparity between races, with the exception that white inventors fare better than all other races. However, for smaller private firms, the discrimination effect reappears, albeit at a less significant magnitude. Finally, for individual inventors, Black and Hispanic inventors are at even more of an unfair disadvantage than indicated by the entire sample.

There is roughly no decrease in discrimination over time. While the hypothesis that increasing awareness of racial and gendered discrimination over time has led to reductions in harm is often somewhat true, it is critical to note that quality-controlled disparities in patenting persist over time. Thus, policy and awareness action are still necessary to bridge the gap.

Further, we apply the same techniques with regard to gender. Male inventors experience an advantage of 1.8% controlling for quality, which is somewhat lower than, for example, the gap between white and black inventors. Nevertheless, unlike race, the gendered gap does not disappear for inventors associated with public firms. This potentially indicates that biased examiners' subconscious distrust of female inventors is stronger than their distrust of Black or Hispanic inventors.

To further confirm causality, we perform human perception tests on the names of inventors. While the USPTO has accurate gender markers produced by machine learning, examiners see only the names. We exploit the fact that many names are of ambiguous gender to show that examiners are determining gender from names and discriminating based on inferred gender; disparities between male and female inventors are much larger among those with obviously gendered names, including when controlling for quality. When feasible, this technique significantly reduces the identifying assumption because whether or not a name is obviously gendered is an intrinsic property of the name and not related to any other aspect of the application.

Interestingly, discrimination effects exist with regard to firm affiliation (indirectly,

socioeconomic status). Specifically, we find significant unfair discrimination against individuals and preferential bias for public firms. In fact, the effect is larger than the demographic dimensions, reaching 6.6 percentage points. While there are many arenas in which affiliation with a large firm ought to make an individual more credible, patents are supposedly judged solely on their content. There is no legal justification for why the same application filed by Apple instead of an individual ought to be accepted instead of rejected.[3] In fact, there exists a compelling argument for preferential treatment towards individual inventors to better encourage new entry into innovation by broader subsets of the population. A decent patent filed by an individual working in their garage ought to be much more promising than the same application filed by Microsoft's patent machine. The existence of the opposite preference suggests patent examiners exhibit a subconscious bias towards "big names," similar to that for males and whites.

Next, we conduct unique and novel tests for intersectional effects between race and gender. The interactions are quite nuanced: for example, among females, black inventors experience somewhat less discrimination, while API inventors experience more of an advantage. Among white inventors, males experience a much larger advantage than in black or API inventors. One of the most important takeaways is the lack of gendered discrimination in API inventors: because these names are less obviously gendered to native English speakers, less discrimination occurs. This strongly suggests that the discrimination effect is purely name-based bias and not driven by confounding factors, resolving one of the most persistent issues with studies of discrimination.

Finally, we also document several additional novel dimensions of discrimination. In addition to proxying for race and gender, names also lead to discrimination on the basis of frequency; inventors with rare names are at a substantial disadvantage. Moreover,

[3]It may be true that the patent would generate more economic value in the hands of Apple due to greater commercialization potential. However, this is irrelevant for the purposes of patent approval. Under 35 U.S.C. § 101, patents must satisfy the criteria of originality, non-triviality, and utility. Thus, the quality and content of the invention itself ought to be the sole determinant of patent acceptance.

inventors located in big cities experience a large advantage. While the average quality is indeed higher in big cities than rural areas, examiners overweight the importance of this distinction, giving an unfair advantage to better located inventors.

A simple homophily effect does not explain many of the results presented in this paper. One interpretation is that examiners intentionally discriminate; we instead show the potential for unfair biases to arise unintentionally as a result of correlation neglect. Because race, gender, and other factors are correlated with quality, examiners mistakenly interpret inventor characteristics as an orthogonal source of information when the correlation is, in fact, already reflected in the content of the patent. Thus, a rational examiner (that does not hold exogenous racial or gendered biases) may make unfairly discriminatory decisions. We provide a theoretical framework of correlation neglect that can lead to discrimination and biases in patenting in Section 3. While this may not be the only source of discrimination, our model provides a unifying framework that can explain all of the dimensions of discrimination and biases documented in our paper.

This paper makes several key contributions. We are the first to simultaneously investigate discrimination across multiple dimensions, examining their relative scale and interactions/intersectionality. We also identify a novel "affiliation bias" in innovation unfairly favoring large and private firms, to an extent greater than gender and racial discrimination. Moreover, we present an alternative "correlation neglect" explanation of discrimination broader than mechanisms proposed by previous research and able to explain the novel bias factors we identify such as inventor location. Finally, we present robust causal evidence that names drive the discrimination effect via human perception tests in conjunction with Causal ML. This methodology has the potential for a wide range of applications throughout the study of discrimination and bias.

The remainder of the paper is organized as follows: Section 2 discusses several areas of relevant literature. A theoretical framework for understanding bias based on correlation

neglect is developed in Section 3. Section 4 describes our techniques for identifying causality and comparisons to earlier methods. Section 5 outlines the data sources and variable used in empirical tests. Causal evidence of discrimination is presented in Section 6, while novel results and interactions are shown in Section 7. Section 8 concludes.

2. Related Literature

This study relates to several strands of research. The first is the general study of economic discrimination against minority groups. For example, Bertrand and Mullainathan (2004), Chetty et al. (2014), Chetty et al. (2020), and Derenoncourt and Montialoux (2021) consider racial disparities in social mobility. Bertrand, Kamenica, and Pan (2015) and Blau and Kahn (2017) examine gender gaps in economic outcomes, mostly due to occupational and industry effects. Hsieh et al. (2019) demonstrates the economic inefficiency arising from the misallocation of talent as a result of discrimination.

Our paper speaks to the study of racial disparities throughout innovation, entrepreneurship, and business. Aghion et al. (2017) and Bell et al. (2019) study how different factors such as neighborhoods "nurture" inventors—such effects could explain the gap in patenting quality we observe. In a similar vein to inventorship, discrimination against minorities has been especially prominent in early-stage start-ups, venture capital, and loans (Chernenko and Scharfstein, 2022; Fairlie, Robb, and Robinson, 2022; Fairlie and Robinson, 2023; Wang, Wu, and Hitt, 2023). Interestingly, discrimination persists beyond entrepreneurship; Cassel, Lerner, and Yimfor (2022) find that Black and Hispanic fund managers experience significant disadvantages. We show yet another avenue of racial discrimination.

The second strand is the analysis of disparities in patenting. From the race perspective, Tabakovic and Wollmann (2018) and Coluccia, Dossi, and Ottinger (2023) both show potential racial discrimination. Meanwhile, Hunt et al. (2012) attempts to explain why the majority of inventors are male in the first place. In terms of ultimate outcomes, Li

and Liu (2023) demonstrates an unfair gendered gap in patent citations while Schuster, Marcowitz-Bitton, and Gerhardt (2022) observe a lack of female participation in transfers from academia to industry and Jensen, Kovács, and Sorenson (2018) generally confirms worse proportional outcomes for women inventors. Few papers unify these two directions. We seek to build upon Cook and Kongcharoen (2010) and Cook, Gerson, and Kuan (2022), who broadly investigate the "idea gap in pink and black," and provide a rigorous, extensive comparison and investigation of interactive effects.[4]

Third, our paper adds to the study of how names relate to economic outcomes. Fryer and Levitt (2004) show a potential alternative, non-causal explanation for the relation between names and socioeconomic outcomes. In the specific case of patenting, we find that distinctively racial or gendered names do lead to discrimination by patent examiners. This is consistent with a wealth of literature that shows the importance of race when an application undergoes review in the job market, including Bertrand and Mullainathan (2004).[5]

Fourth, our paper complements an extensive corpus of literature studying the unreliability of patent examiners. For example, Frakes and Wasserman (2017) shows the variance in standards based on time allotted to examiners. Many studies apply examiner assignment as an instrument, such as Farre-Mensa, Hegde, and Ljungqvist (2020) and Sampat and Williams (2019). Yang (2024) argues that examiners may prioritize applications with strong writing despite weak underlying technology. We provide strong suggestive evidence that investors discriminate solely based on inventor names.

Finally, this paper contributes to the application of machine learning to study bias and causal effects. Athey and Imbens (2015) provide one of the first causal applications of machine learning. Chernozhukov, Newey, and Singh (2022) investigates the need for debiasing when applying causal machine learning techniques. Specifically in terms

[4]Another interesting direction is mentioned in Lerner (2002), who observe that some countries historically prefer domestic over foreign applicants.

[5]See Bertrand and Duflo (2017) for an extensive discussion of such experiments.

of textual analysis, both Veitch, Sridhar, and Blei (2020) and Khetan et al. (2022) study embedding based techniques for causal inference. Hochberg et al. (2023) use similar techniques to demonstrate that female inventors receive less citations. We use more powerful OpenAI embeddings to study and compare patenting discrimination across dimensions.[6]

3. Theoretical Framework

Why are there unfair biases in patent approval? One possibility is that the examiners hold intrinsic personal biases against certain groups and intentionally discriminate. However, the disparities we document in this paper could also arise unintentionally. In this section, we model the decision-making of a patent examiner in order to demonstrate a correlation neglect explanation for unintentional discrimination. Broadly, examiners may overweight factors such as firm affiliation or race because they mistakenly interpret them as independent signals of innovative strength.

Specifically, the role of the patent examiner is to determine the quality q of a patent application. Three factors may drive the quality of an invention: first, the ingenuity of the idea itself, q_i. Second, the background and upbringing of the inventor, or q_u. Akcigit, Grigsby, and Nicholas (2017), Bell et al. (2019), and others document the importance of environmental factors such as parental income or childhood exposure to entrepreneurship in nurturing innovative skills. We deem these factors unobservable because they are not included in the patent application. Third, the current resources available to the inventor, q_o. An inventor affiliated with Apple and with access to a successful patent lawyer will, on average, produce higher-quality inventions. Such information is often provided throughout the patent application, and thus, this factor is observable to the examiner. We

[6] Chernenko and Scharfstein (2023) cautions against the use of certain race-identifying algorithms when studying racial bias. See Section 5.2 for a discussion of why such techniques are still applicable in cases such as this paper.

let

$$q = q_i + q_u + q_o. \tag{1}$$

In order to estimate q, the examiner can read the content of the application, which provides a noisy signal $s_1 = q + \epsilon_q$. If the examiner is perfectly skilled, $\epsilon_q = 0$. In practice, this is never the case, and thus the examiner seeks other signals. By seeing the firm (or lack thereof) the inventor is affiliated with, the city the inventor resides in, and other information, the examiner can infer the resources and social network of the inventor, $s_2 = q_o + \epsilon_o$. Finally, the background of the inventor is not directly observable. Nevertheless, socioeconomic, exposure to innovation, and other conditions are highly correlated with demographic characteristics such as race and gender. Because names are often correlated with race and gender, the examiner obtains a (very noisy) signal $s_3 = q_u + \epsilon_u$ from the inventor's name. We assume all variables $(q_i, q_u, q_o, \epsilon_q, \epsilon_u, \epsilon_o)$ are independent and normally distributed with zero mean and variance $(\sigma_i^2, \sigma_u^2, \sigma_o^2, \sigma_{\epsilon_q}^2, \sigma_{\epsilon_u}^2, \sigma_{\epsilon_o}^2)$. The examiner then determines $E[q|s_1, s_2, s_3]$.

Proposition 1. *Given signals s_1, s_2, and s_3, the conditional expectation of patent quality q is given by*

$$E[q|s_1, s_2, s_3] = w_1 s_1 + w_2 s_2 + w_3 s_3 \tag{2}$$

where

$$w_1 = \frac{\sigma^2}{\sigma^2 + \sigma_{\epsilon_q}^2}, \quad \sigma^2 = \sigma_i^2 + \frac{\sigma_o^2 \sigma_{\epsilon_o}^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} + \frac{\sigma_u^2 \sigma_{\epsilon_u}^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2},$$
$$w_2 = \frac{\sigma_o^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} \cdot (1 - w_1), \quad w_3 = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2} \cdot (1 - w_1).$$

The proofs of this and later propositions are provided in the Appendix. Importantly, if σ_{ϵ_q} is small relative to σ_i^2, then w_1 will approach 1. Essentially, if the examiner is good and the innate quality of the idea is important, then the examiner should only use the text of the application in their estimation of q, i.e., $E[q|s_1, s_2, s_3] \approx s_1$.

However, the examiner may overlook the correlation between the signals s_1, s_2, s_3. This is a well-documented effect known as *correlation neglect*, where agents do not understand that distinct signals contain repetitive information. Enke and Zimmermann (2019), Eyster and Weizsacker (2011), and DeMarzo, Vayanos, and Zwiebel (2003) study the causes and effects of similar biases in various settings.[7] In the case of the patent examiner, one may be incognizant of the fact that s_1 already captures all three factors and instead mistakenly interpret the three signals as $s_1' = q_i + \epsilon_i$, $s_2 = q_o + \epsilon_o$, $s_3 = q_u + \epsilon_u$. Essentially, the examiner thinks that they are gleaning additional information about quality from the affiliation and name of the inventor even though it is already contained in the actual content of the invention.

Proposition 2. *1) When the examiner misinterprets s_1 as s_1' under correlation neglect assumptions, the conditional expectation of patent quality q is given by*

$$E[q|s_1', s_2, s_3] = w_1' s_1' + w_2' s_2 + w_3' s_3 \tag{3}$$

where

$$w_1' = \frac{\sigma_i^2}{\sigma_i^2 + \sigma_{\epsilon_q}^2}, \quad w_2' = \frac{\sigma_o^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2}, \quad w_3' = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2}.$$

2) The following always holds:

$$w_1' < w_1, \quad w_2' > w_2, \quad w_3' > w_3.$$

The correlation neglect of the examiner results in an overestimation of the importance of s_2 and s_3, i.e., *affiliation bias* and gendered and racial discrimination. Because the optimal strategy for a skilled examiner is essentially to use only the textual signal s_1, a double-blind

[7] Interestingly, Levy and Razin (2015) show potential positive effects of correlation neglect in voting behavior. Because our examiners are otherwise rational agents, this bias is inefficient here.

reviewing process would significantly reduce the risk of bias while having minimal impact on approved application quality.

4. Causally Identifying Discrimination

4.1. An AI-Quality Approach

The biggest challenge in studying bias lies in distinguishing racial disparities driven by structural inequities and other outside factors from unfair discrimination on the part of the regulator. While both present significant social challenges, they necessitate significantly different solutions. Isolating these two effects can be very challenging since it requires the determination of the merit of an application.

In the case of patent applications, all of the legally necessary information to determine worthiness lies in the text of the application. Thus, we seek to construct an objective measure of patent quality based on the contents of the application text. With such a measure as a control, we may effectively test a natural experiment by observing the outcomes of equivalent applications (in quality) filed by inventors from different demographic groups. This approach has the identifying assumption that the measure must accurately determine quality and that the text contains all relevant information.[8]

The textual information in a patent application is very specific and high-dimensional, so we take an embedding-based approach to capturing quality. Specifically, we construct an embedding-based measure similar to that presented in Yang (2024).We first take a Large-Language Model (LLM) embedding (semantic vectorization of a text) of the title and abstract of each patent using OpenAI's Ada-002 model. This captures the meaning of the entire body of the text into a $1,536$-dimensional numerical vector. LLM embeddings are very powerful and complete semantic representations of texts in their entirety (Rad-

[8]By definition, the quality of a patent application is a function of its contents, so in the case of patenting, the second assumption is always satisfied.

ford et al., 2018; Brown et al., 2020). In the context of patent applications, they capture information related to a variety of patent outcomes from economic value to citations, encapsulating the essence and quality of the invention (Yang, 2024). Then, we use this embedding vector as a primary input along with other fundamental variables (Cooperative Patent Classification and the number of claims) to secondary machine learning models, identifying whether ann application deserves to be accepted. Specifically, the embedding is fed into a three-layer neural network with activation function *Mish* as proposed by Misra (2019):

$$Mish(x) = x \cdot \tanh(\ln(1 + e^x)). \tag{4}$$

The models predict the binary outcome *Acceptance*, equal to 1 if an application is accepted and 0 otherwise. The net acceptance rate of applications is over 70%, so we optimize the neural network for binary cross-entropy to account for this skewness. Moreover, the models are trained with a three-year rolling window sample, i.e., the predictive model for 2009 is trained on data from 2006 to 2008, ensuring all predictions are out-of-sample. Generally, the predictions are very accurate, with F1 Scores of over 80% across the board. We define *application quality* as an application's predicted chance of acceptance, i.e., how much its contents "deserve" acceptance.

Essentially, the models are trained to take as input the text of the application (i.e., description of the invention) and determine objective *application quality* from its contents. With this measure as a control, we can identify unfair outcomes in the aggregate, again assuming that the model captures textual quality sufficiently.

4.2. Human Perception

While the ML quality measure is very accurate, it is impossible for it to be ideally so. In general, Causal ML provides strong evidence of causality but does not absolutely rule out

confounding factors. We propose a novel parallel methodology relying on human perception to resolve this concern. Ultimately, the goal of identification in the case of patenting, job applications, loans, and many others is to determine whether the decision-maker is using name/demographic information, for to do so is illegal. Causal ML techniques, such as our quality measure, demonstrate an even stronger result, namely that the discrimination is also unfair. At the same time, we now present a parallel method that can produce even more decisive proof of the illegal use of name information.

Crucially, actual demographic information is not available to patent examiners at the time of review because of discrimination concerns. Racial and gendered discrimination then happens, if at all, through the examiners' perception of names. Take the example of bias against female inventors. If an inventor is female but has a name that is challenging to identify as female, the examiner would likely not discriminate. Thus, we manually identify through inspection the "inferred gender" of inventors through their reported names. Then, we exploit the difference between this inferred gender and the "real" gender (generated by the USPTO) to reject the alternative hypothesis that inventors do not use names, and the Causal ML results result from a correlation between gender and quality uncaptured by the quality measure. Specifically, if inventors are not inferring gender from names, it should make no difference whether the name is easily gendered, and only the true gender should be related to outcomes. Thus, if whether a name is gendered is correlated to outcomes, the examiner must be illegally using name-based information in some way.

This technique can be applied in conjunction with a Causal ML control such as our *application quality* measure. In the case of gender, one may perform two quality-controlled regressions separately on a sample of genderable and gender-ambiguous names. This substantially reduces the identifying assumption. The only case where a spurious correlation can occur is if the "genderability" of a name is significantly correlated with a dimension of quality not captured by the ML measure. The only assumption needed is then the negation

of this scenario, which is almost certain to be true.

The AI + human perception approach presents significant opportunities in various fields, particularly in identifying biases arising from the judgment of human decision-making, such as in job or loan applications. In contrast to other variables, the impression of a name should have no correlation to the actual content of the application, and its correlation with outcomes demonstrates that the decision-maker is illegally using the name.

4.3. Existing Methodologies

The most traditional approach for rigorously determining the presence of discrimination is performing a true field experiment as in Bertrand and Mullainathan (2004). Because the information available to (unknowing) participants in field experiments can be wholly controlled, it can be very clearly determined that, for example, hiring managers are making decisions based on the name of the applicant. This is not feasible in the case of patenting, as discussed in the introduction.

One alternative approach developed in the computer science literature uses textual analysis techniques to identify treatment effects. Hochberg et al. (2023) apply this technique and discover a relation between gender and patent citations. Essentially, one trains separate models to simulate outcomes exclusively for each gender, then cross-applies the models on the exact text to approximate counterfactual treatment effects. To identify causality, this technique requires the same identifying assumptions as our ML Quality approach, namely that the text contains all relevant information and that the ML model captures all textual information. Otherwise, the result could be a result of a confounding factor (such as upbringing) driving a correlation between demographic groups and outcomes.

However, such a method also presents several challenges resolved by our unconditional quality approach. First, training counterfactual models becomes prohibitively expensive and difficult to interpret when discrete variables with many possible values (such as

race) and interaction effects are involved. For example, the interaction between race and affiliation would produce fifteen distinct groups, each necessitating its own model to then be cross-applied to all other groups.[9] Thus, this technique is usually restricted to determining a binary treatment effect, whereas we apply our quality control to a wide range of possible bias factors. Further, our methodology also allows for determining continuous effects, such as the name frequency results in Section 7.4.

5. Data and Variables

5.1. Data

Our primary data source is the US Patent and Trademark Office's (USPTO) PatentsView service, which provides bulk textual data for all utility patents granted after 1976 and all patent applications published after 2001. Firm market capitalization is obtained from CRSP.

5.2. Variables

AI Quality is produced as described above in 4.1. Our measure of individual inventor race is constructed using the Bayesian Improved First Name Surname Geocoding (BIFSG) algorithm via the Surgeo Python package. We note that Chernenko and Scharfstein (2023) recently found that the use of such algorithms may lead to substantially biased results. However, this method is appropriate in our use case for two reasons: first, we wish to create a measure of "perceived" race, in order to measure the potential for discrimination. Thus, even if the algorithm is inaccurate, our results will be accurate as long as the race prediction matches the patent examiner's perception of race. This is the case because patent examiners do not interact directly with or observe a self-reported race from the patent

[9] In cases of such interactions, some groups alone may not contain enough observations to train a distinct, sufficiently accurate ML model in the first place.

inventors, and thus any bias is introduced from the "race" of the name. Second, much of the potential error in empirical tests using these algorithms stems from the geographical element. However, for our primary tests, we set the location for all inventors to be Alexandria, Virginia, the location of the USPTO headquarters and examination center. For robustness, we also test without geocoding, and obtain largely the same result. Geocoding for inventor location is unreasonable since the examiner's subconscious evaluation of race would not take into account the racial distribution of, say, zip codes in Seattle. The measure of gender is created by the USPTO and taken directly from PatentsView, which uses machine learning algorithms to identify the most likely inventor's gender based on first and last name. There are no self-reported demographics available in the bulk patent data.

For most tests, race is measured using a dummy variable measuring the "plurality" race of a patent. For example, if a patent has five listed inventors, three Hispanic, one white, and one Black, then the plurality race is Hispanic. This same logic applies to individual inventors, as SURGEO outputs continuous probabilities for the race and not a discrete binary value. That is, an individual is assigned the race to which they most likely belong.

6. Racial and Gendered Disparities

6.1. Suggestive Evidence of the Racial Gap

The motivation for this study lies in the significant differences found in innovation outcomes across races. Specifically, while a wealth of research studies later-stage outcomes, such as VC funding (Cook, Marx, and Yimfor, 2022) and citations (Hochberg et al., 2023), the area of patent applications remains fundamentally understudied. This is puzzling when considering the importance of patent successes especially in the survival of young firms and startups, and potential implications for social mobility. We hypothesize that patent applications filed by white inventors will, in aggregate, be more likely to be accepted,

whereas those filed by minority groups may be less successful.

Table 1 reports the summary statistics of application acceptance and predicted acceptance rate as sorted by race. Here, race at the patent level is defined as the plurality race when averaging inventors together as a group. In the Appendix, we conduct robustness tests for different definitions of patent-level race. API patents are accepted the most, with an acceptance rate of over 70% across all years, while white inventors receive the second highest rates, at 67.7%. Black inventors receive the worst outcomes of all groups except Indigenous Americans, with only a 60.6% acceptance rate. Comparing instead with predicted acceptance rates in Panel B, the relative rankings of groups by mean rate remain the same. However, the mean predicted acceptance rate for black inventors is 66.8%, 2% higher than the actual outcomes, whereas both White and API are much closer in terms of difference. These statistics seem to suggest that race may be influencing patent acceptances since the model's predicted rates do not entirely capture the disparity between groups.

[Insert Table 1 Here]

6.2. Identifying Mechanisms

To more rigorously determine this potential effect, we first turn to a standard regression setting. By regressing patent acceptances on race indicator dummies, we are able to quantify the role that race plays. Crucially, we are able to control for the objective measure of application quality provided by the predicted acceptance rate from the machine learning model. Specifically, we estimate the following regression for application-time (i, t):

$$Accepted_{i,t} = \sum_{R \in \{Races\}} \beta_n \cdot R_{i,t} + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{5}$$

Due to collinearity, the set $\{Races\}$ cannot contain all six race variables. In general, Hispanic inventors as a group tend to roughly rank in the middle of the races, in terms of

group size and also potential discrimination (i.e., the difference between predicted rates and real acceptances). Thus, the Hispanic group is treated as the benchmark and reported coefficients from the regression are to be viewed as relative to the Hispanic group. That is, if being White grants an unfair advantage to inventors, this is relative to Hispanic inventors.

Table 2 reports these regression results. The relation between race and patent acceptance remains significant at the 1% level under all specifications for the groups White, Black, and API. Black inventors are systematically discriminated against, while White and API inventors receive an advantage. Importantly, these coefficients do drop by half or more when introducing the application quality control. Thus, the disparity in overall success is driven by two orthogonal, roughly equally important factors: first, Black and Hispanic inventors, on average, produce lower "quality" inventions than White and API inventors (where quality is determined as per the USPTO's metrics for patent acceptance), and that, *ceteris paribus*, they still experience lower acceptance rates.

[Insert Table 2 Here]

Overall, controlling for application quality removes almost all other confounding factors that could be driving a spurious relation. For example, if inventors were disproportionately sorted into industries with different aggregate acceptance rates by race, then summary statistics would show a disparity in successful innovation outcomes. This effect is controlled for because the application quality measure accounts for sectors and matches the aggregate acceptance rates across industries (see Yang, 2024, Figure 1). Of course, there may still be confounding factors driving this relation. One primary source of variation is found in patent examiners, who may be individually harsher, more lenient, or busier (Farre-Mensa, Hegde, and Ljungqvist, 2020; Shu, Tian, and Zhan, 2022). If white inventors were to receive more lenient examiners, this could also drive the results. However, this scenario would still qualify as unfair discrimination, and the USPTO's claim of true random

assignment is generally accepted in the literature.

6.3. Cross-Sectional Analysis of Racial Disparities

The most obvious explanation for the observed unfair treatment is racial discrimination, intentional or subconscious, on the part of the examiner. One soft, unregulated factor that may influence examiner decisions is individual characteristics of the inventor. Because examiners work with deadlines and cannot expend too much effort evaluating one application, first impressions likely play a large part in the ultimate decisions. An examiner may immediately judge a patent as acceptable for several reasons. For example, if the writing and other portions of the application are professional and strictly adherent to guidelines, an experienced lawyer may have worked on the patent. Second, if the patent is filed by a large firm such as Apple or Microsoft, it likely contains novel, useful technology. However, it is also possible that certain examiners may view inventors with generically African American names as lower quality because of their innate personal biases, and vice-versa for white-sounding names. Under correlation neglect, the results for weak applications filed by minorities would be skewed downwards.

If this is the case, the disparities between races should be larger for individual inventors and smaller applicant firms and smaller (or non-existent) for applications filed by public firms, where a big name may be powerful enough to push the application over the line alone. To test this hypothesis, we split the sample into three cross-sections: public firms, private firms, and individuals. Public firms are all assignees that can be matched to a publicly traded PERMNO in CRSP. Private firms are assignees that are not matched to public companies or government agencies. Individuals are simply inventors filing for a patent under their name.

6.3.1 *Public Firms*

In Table 3, we report, for the subsample of publicly traded firms, the same summary statistics and regressions as for the full sample. Interestingly, many of the trends seen in the full sample do not apply for public firms. Across races, there is much less variation, both in acceptance rate and application quality.[10] In the regression setting, white inventors still experience some advantages, but the effects are much smaller and only appear in a few specifications.[11] Clearly, the discrimination effect is much less prominent for applications filed by public firms.

[Insert Table 3 Here]

6.3.2 *Private Firms*

Table 4 shows the same tables for the subset of applications filed by private firms. These represent a middle ground between big public firms and individual applicants—an LLC or Corp. will reassure the examiner, but not to the level that NVIDIA or AMD would. Here, although the acceptance rates look somewhat similar across races, white inventors experience a significant unfair advantage when controlling for quality. Overall, the discrimination effect is much stronger than found in public firms, but still weaker than in the entire sample, especially when comparing the relative pairs of White-Black and API-Black.

[Insert Table 4 Here]

6.3.3 *Individual Inventors*

Individuals are the most likely to experience an unfair process in patenting: for example, many inventors may lack a patent lawyer. They are also, unfortunately, the most vulnerable.

[10] The closing of the gap in quality could be a result of several factors. For example, greater access to patent lawyers when working under a major company, or talent identification in the hiring process.

[11] The observation count of inventors identified as Indigenous American by SURGEO is too low to draw conclusive results from, but the evidence generally suggests a moderate level of discrimination.

While an inventor at Apple is unlikely to lose their job simply because of a single failed patent application, an aspiring entrepreneur may find themselves in financial ruin after unsuccessfully trying to protect their intellectual property. If individuals are discriminated against on the basis of race, this would further exacerbate the problem and have significant implications for socioeconomic mobility.

Table 5 reports the results for individuals. For individuals, even the summary statistics show that acceptance rates are not at all close across races. While Black inventors experience an average rate of 54.3%, 60% of White patents are granted, a difference of 5.7 percentage points. While there is a difference in quality between these two groups, it is significantly lower, at only 2.2 percentage points. When controlling for application quality, white and API inventors experience a major advantage relative to the Hispanic group, and the Black inventors have a major disadvantage relative to Hispanic inventors and, of course, White/API. The magnitude of the effect is even greater than in the full sample. The true extent of the discrimination effect is, in fact, masked by the presence of public firms in the sample.

[Insert Table 5 Here]

6.4. The Gendered Patenting Gap

In addition to race, another important dimension of discrimination is gendered bias (essentially, misogyny). Like race, gender is not officially reported in the patent application process, but the name of the inventor can allow the examiner to (either subconsciously or consciously) form their own determination of the inventor's likely gender. We define gender as the majority gender among inventors of a certain application.

Similar to race, disparities exist both in application quality and in aggregate acceptance rates. Table 6 reports these summary statistics. Specifically, the mean acceptance rate for female majority patents is only 63.8%, whereas the male acceptance rate is 69.0%. The gap

in quality (predicted rate from text) is smaller, from 65.2% to 69.6%.[12]

[Include Table 6 Here]

Next, we perform the same quality-controlled test as in Section 6.2. Specifically, we estimate the following regression:

$$Accepted_{i,t} = \beta \cdot Male_{i,t} + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}, \tag{6}$$

where Male is a binary variable equal to 1 if the inventors are majority male and 0 otherwise.

Table 7 reports the results. Across all specifications, applications filed by male inventors are much more likely to be accepted. Controlling for quality and year-fixed effects, the effect is still nearly 2 percentage points. From a magnitude perspective, this bias is stronger than the gap between Hispanic and Black or Hispanic and white but slightly weaker than the gap between Black and white. Clearly, both issues are of great importance, and similar policy actions may be necessary in both areas. For example, the simplest remedy would be to remove inventor names from the data provided to the examiner and instead only record the names for ex-post analysis and actual assignment of IP rights.

[Include Table 7 Here]

6.5. Cross-Sectional Analysis of Gendered Disparities

The natural question that arises for gendered discrimination is what results the cross-sectional analysis by applicant type presented in Section 6.3 would produce. Ex-ante, the result is unclear: the backing of a large firm could have an ameliorating effect on the bias, or the discrimination could persist across all applicant types. Thus, we conduct the same

[12]Note again that quality accounts for many dimensions of variation in acceptance rate, including field, and is not necessarily indicative of inventor skill. For example, patent applications filed in a more competitive field are, by definition, lower "quality."

quality-controlled regression of application acceptance and gender for all three applicant types (Large Firms, Private Firms, and Individual Inventors).

Table 8 reports these regression results. Interestingly, the effect remains large and statistically significant in all three categories. While there is a slightly larger discrimination effect for individual female inventors, the magnitude of the difference is negligible. Essentially, female inventors experience a substantial disadvantage *irrespective* of their affiliation with a large firm. This is conceptually consistent with previous results showing that female employees at innovating firms experience unfair treatment throughout the innovation chain and not only in patenting (Waldfogel, 2023; Chien and Grennan, 2024).

Note that only a simultaneous analysis of gender and race with identical techniques as presented in this paper can demonstrate and quantify differences between gendered and racial discrimination. There are several possible explanations for this effect in particular, with the most obvious being that the gender bias is "stronger." In the next section, we investigate potential explanations for this difference and compare various dimensions of discrimination.

[Include Table 8]

6.6. Gendered Discrimination via Human Perception

The aforementioned substantially reduced effect of gendered discrimination in API inventors provides some evidence that examiners are directly inferring gender from names. To further demonstrate this mechanism, we exploit the ability of ML algorithms to label gender beyond human inference. Specifically, the USPTO provided gender variable uses ML to determine inventor gender accurately.[13] Meanwhile, we now manually label the $15,000$ most common inventor names as obviously male, obviously female, or not obviously gendered. When manual labeling is successful, it almost always matches

[13]This measure is generated by the USPTO for ex-post statistical studies and is not available to the examiner at the time of review.

with the ML measure, but many ML-labeled male or female inventors are not obviously gendered to a human reader.[14] This difference provides the statistical power for our test: if examiners are discriminating based on reported inventor names, the effect should be much weaker for inventors with names manually labeled as not obviously gendered.

Specifically, we split the sample into two parts. First, inventors with names obviously gendered to the average American English-speaking reader. Second, inventors without obviously gendered names. Then, we perform the same gendered discrimination regression as in Section 6.4.

Table 13 reports the results. Although males experience an advantage in both samples, the effect is universally twice as large in inventors with obviously gendered names.[15] This difference is strong evidence that examiners are not examining other aspects of the application correlated with gender but instead are literally making judgments based on an impression of gender derived from the inventor's name. Switching to a double-blind procedure could thus cut unfair discrimination against female inventors at least in half, if not more.

[Insert Table 13 Here]

7. Discrimination and Heterogeneity Across Patent Applicants

7.1. Applicant Types and Affiliation Bias

Essentially, the cross-sectional tests presented in Section 6 check for interactions between applicant type and race/gender. However, why should applicant type only be used as a filter to better understand racial and gendered discrimination? In fact, why would

[14] In untabulated results, the manually-labeled gender measure produces slightly stronger results than the USPTO gender as a regressor, confirming the labeling accuracy.

[15] In order to not over-estimate the difference between samples, we lean towards ambiguity when labeling names. Thus, some names that may seem female or male to certain examiners are labeled ambiguous, which could be the cause for the persistence of the effect in ambiguously gendered names. Regardless, our result effectively provides a lower bound on the actual impact.

patent examiners (already shown to be unfair) limit their subconscious biases to the most obvious demographic groups? A rational examiner experiencing correlation neglect would use all information available. Thus, discrimination across non-traditional dimensions such as applicant type could also be economically important.

For example, a patent examiner may subconsciously give more credence to claims made by inventors from Apple or Microsoft than John Smith from Arkansas, even if the two file exactly identical applications.[16] This can potentially be incredibly harmful in terms of social mobility since poorer entrepreneurs and inventors will obviously file applications as individuals.

Table 9 reports the aggregate acceptance rates and mean application quality across applicant types. Notably, the differences here are much larger than in gender and race, while the sizes are as expected: large firms have the lion's share of their applications accepted, at 75.1%, with far less accepted for private firms and individuals (66.7% and 60.1%). Again, differences in quality/expected acceptance also exist, but are notably several percentage points lower than the gaps in actualized outcomes. In the case of applicant types, it is expected that there will be large gaps in quality irrespective of unfair treatment; after all, if NVIDIA were producing worse patent applications than individuals, that would be cause for significant concern.

Following these results, we estimate the following quality-controlled regression:

$$\textit{Accepted}_{i,t} = \beta_1 \cdot Public_{i,t} + \beta_2 \cdot Private_{i,t} + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{7}$$

The regression results are shown in Table 10. In all specifications, private firms experience an advantage relative to individuals, and large firms an even larger advantage. In

[16] Whether such a bias is efficient is not immediately clear, but it seems rather inappropriate in the context of patent applications. Unlike, for example, credit score in loan or mortgage applications, the identity of the inventor should be irrelevant in such a technical context.

fact, the gap between large and private firms is twice as large as that between private firms and individuals. Overall, large firms are associated with an improvement of almost 7 percentage points in acceptance rate *with the same quality of application*. The persistence of such a large effect even controlling for quality is concerning because, again, the ability of individuals to succeed and break through in innovation is critical to entrepreneurship and social mobility.

[Insert Table 10 Here]

More broadly, a homophily-based theory of discrimination/bias is not able to explain the discrimination effect for applicant types.[17] Instead, it seems that, likely based on socialization, examiners have implicit biases towards viewing certain individuals as more credible and others as less. A strictly homophily based view of the world may, in fact, restrict our understanding of discrimination given that many dimensions of bias may exist based on non-demographic factors. While homophily effects certainly exist, they are not the only, and in this case, not the most important, explanation for bias.

7.2. Heterogenous Racial and Gendered Discrimination

If examiners are discriminating, would the effects be homogenous? One strong hypothesis is that disadvantaged inventors would experience more significant penalties when submitting weaker applications since any mistake will be weighed more heavily. For example, an examiner with subconscious racial biases may artificially inflate the magnitude of a mistake made by a black inventor. To test for this effect, we examine the effects of the three dimensions of discrimination (race, gender, and applicant type) in relation to application quality.

[17] Unless one considers the government as a large firm and argues that employees of the government identify themselves implicitly with corporations. However, this seems rather unlikely. If anything, average employees of large firms empirically tend to display dislike towards corporations in general.

Specifically, we separate the sample into below- and above-median quality applications and test for discrimination in both halves of the sample. We consider the most restrictive regression settings in which fixed effects and application quality controls are included, and standard errors are clustered by year. The coefficients and their 95% confidence intervals for the high-quality and low-quality samples are plotted in Figures 1 to 3.

Figure 1 shows that there are substantial differences in racial disparities based on application quality. White and API applicants' advantages are substantially higher for low-quality applications. Racial disparities are substantially lower among high-quality applications, with some statistically insignificant from zero.[18]

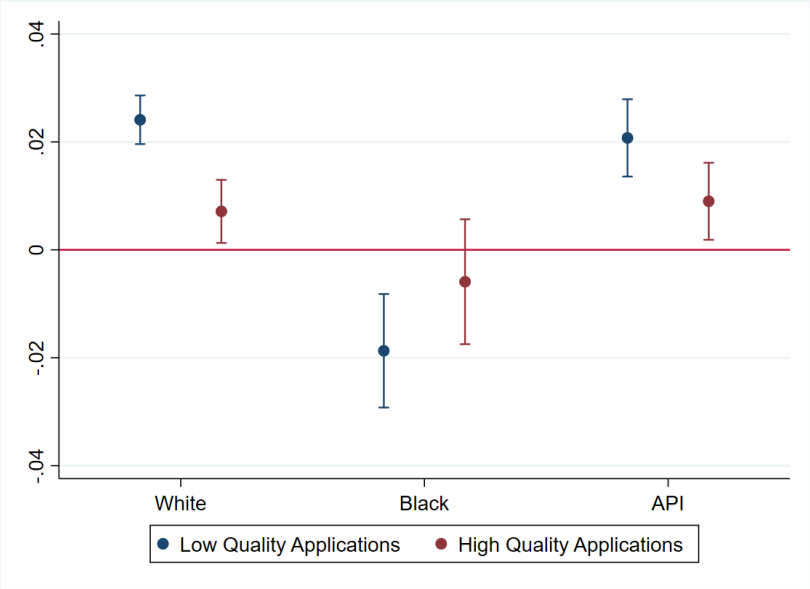


Figure 1. Race Disparities: High- vs Low-Quality Applications

Similarly, Figure 2 reveals that the gender disparity is also significantly greater for low-quality patent applications. Figure 3 finds that large firms enjoy substantial advantages for all applications filed, with a greater advantage for low-quality applications (relative to individual inventors). Private firms only experience an advantage for low-quality applications.

[18]The difference in the proportional effect is, in fact, greater because low-quality applications have a lower

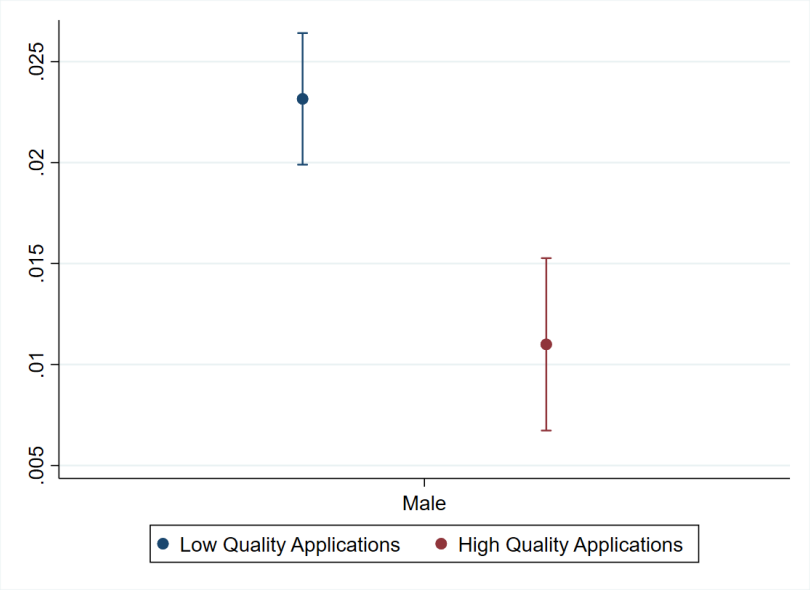


Figure 2. Gender Disparities: High- vs Low-Quality Applications

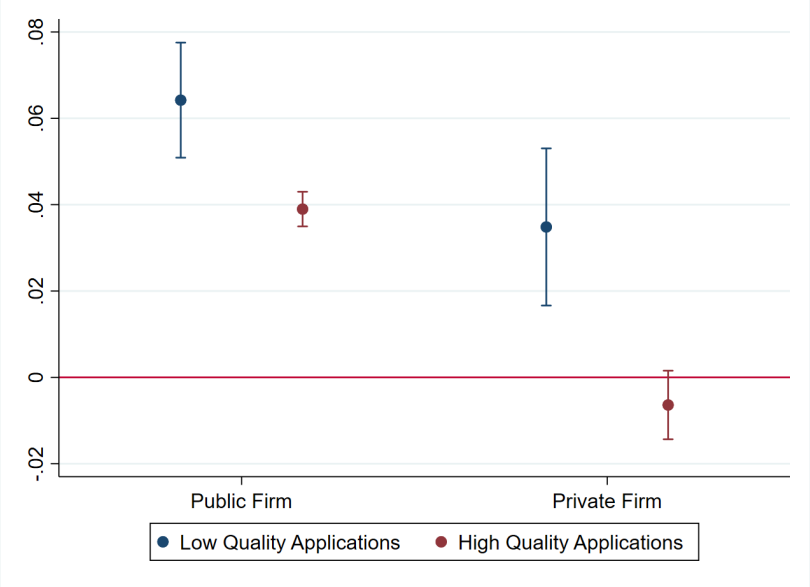


Figure 3. Disparity in Applicant Types: High- vs Low-Quality Applications

predicted acceptance rate.

7.3. Effects of Intersectionality and Bias Interactions

One benefit of applying a uniform methodology to varying dimensions of discrimination is to compare their magnitudes and identify cross-sectional differences. However, we can also investigate the effects of intersectionality between groups. For example, would membership in more than one disadvantaged group ameliorate or compound the discrimination effect? The answer is ex-ante unclear, but it is rather unlikely that the interaction would be simply additive.

Specifically, we are particularly interested in the intersectionality between race and gender, which is an important field. To identify the subtle interactions between the two, we perform two interaction tests. First, we filter on race, then perform gender regressions. Second, we filter by gender and perform race regressions.

Table 11 shows that conditional on inventor gender, racial discrimination persists but with substantially varying effects. The advantage experienced by white and API inventors is greater for female than male inventors, while the disadvantage experienced by Black inventors decreases. In fact, API females enjoy a much more significant advantage (3.5%), which even exceeds that of white females (1.9%). Among males, white and API inventors have roughly equal advantages (1.4%) and Black inventors a disadvantage (-1.4%).

[Insert Table 11 Here]

Table 12 displays the gender disparity conditional on inventor race. Again, there are substantial cross-sectional differences in discrimination. Gender bias is strongest among applications filed by Hispanic (2.7%) and White (2.5%) inventors, followed by Black inventors (1.4%). Notably, gender disparity is almost entirely mitigated for the API population (0.7%). This supports the hypothesis that the discrimination is purely driven by the inferred gender from the name since API names are often harder to accurately identify as male or female for the average American examiner. In the following section, we conduct an extensive test to verify that the name alone drives the bias.

[Insert Table 12 Here]

7.4. Additional Bias Factors

Broadening our examination of discrimination beyond demographic groups and homophily effects also opens the door to several other dimensions of earning credibility in the eye of the examiner. One is affiliation with a powerful institution, as discussed in Section 7.1. However, many other "soft" variables are also available to the examiner. For example, we have already shown that examiners discriminate by race and gender via inventor names. However, there could be a third dimension: bias against uncommon names. For example, an examiner could prefer a patent filed by "John" rather than "Percival," even though both are most common among white men. Because name frequency over the entire population is not representative of inventors in particular, we define Name Frequency as the natural logarithm of the total occurrences of a particular first name over all inventors minus one. Then, we estimate the following regression:

$$Accepted_{i,t} = \beta \cdot Name\ Frequency i, t + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{8}$$

Panel A of Table 14 reports the results. The effect is quite pronounced with the most common names experiencing an advantage of over 6% relative to uncommon names. However, this could also be a result of name frequency simply proxying for male or white inventors, both of which are over-represented. Thus, we perform the same regressions controlling for both race and gender.

Panel B of Table 14 shows the controlled regression. Interestingly, the effect remains as strong or grows even more significant when controlling for race and gender, suggesting that name frequency is an orthogonal dimension of discrimination.

[Insert Table 14 Here]

Another potential source of bias lies in an inventor's location, which is also reported on the patent application. Under the correlation neglect hypothesis, an examiner may (correctly) understand that the average innovation produced in San Francisco is more impactful than that produced in rural areas yet still overweight the importance of the location in their judgement. To test for such an effect, we classify American inventors as either from large cities or not, where a large city is defined as a top 50 most populous city in the United States. Then, we estimate the following regression:

$$Accepted_{i,t} = \beta \cdot Large\ Cityi, t + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{9}$$

Table 15 shows that examiners do favor inventors from big cities, but the effect is relatively small in magnitude when controlling for application quality. Examiners are less biased against inventors from rural areas or smaller cities than against those without firm affiliations. This is consistent with the relatively low importance of inventor location. For example, firm affiliation is a stronger predictor of quality in the first place. Thus, under correlation neglect, examiners will exhibit more significant beneficial treatment to such inventors than those from large cities. Nevertheless, bias along this dimension persists and would presumably also disappear via double-blind reviewing.

[Insert Table 15 Here]

8. Concluding Remarks

We use generative AI and human perception techniques to identify USPTO discrimination against under-represented groups and individual inventors in violation of U.S. Code. This paper is the first to thoroughly investigate discrimination across various dimensions and the interactions/intersectionality among them. We are also the first to identify a

significant "affiliation bias" in innovation unfairly favoring large and private firms to an extent more significant than that of gender and racial discrimination. Bias against individuals based on such affiliation is unlikely to be unique to patenting and thus necessitates further study in a range of fields. We explain all of our results through a novel theoretical framework based on correlation neglect, justifying discrimination by rational regulators in the absence of ill intent, which double-blind practices can resolve.

These results show the persistence of a significant roadblock to achieving socioeconomic mobility and equality. Educational opportunities and other environmental factors limit the development of skills in under-represented and individual inventors, but unfairly discriminatory treatment by the government compounds existing disparities. These two problems necessitate a dual-pronged approach on the part of policymakers, preferably beginning with the immediate reform of USPTO practices to be double-blind.

Meanwhile, the methodology presented in this paper opens opportunities for the comprehensive study of discrimination in any area involving a human decision-maker, applications/requests by individuals, and allocation of a scarce resource. In many cases where field experiments are infeasible, an AI + human perception approach can rigorously distinguish correlations between demographics and outcomes from human bias, which is critical to identifying solutions and rectifying such issues.

References

Aghion, P., U. Akcigit, A. Bergeaud, R. Blundell, and D. Hémous. 2019. Innovation and top income inequality. *The Review of Economic Studies* 86:1–45.

Aghion, P., U. Akcigit, A. Hyytinen, and O. Toivanen. 2017. The social origins of inventors. Working Paper, National Bureau of Economic Research.

Akcigit, U., J. Grigsby, and T. Nicholas. 2017. The rise of american ingenuity: Innovation and inventors of the golden age. Working Paper, National Bureau of Economic Research.

Altonji, J. G., and R. M. Blank. 1999. Race and gender in the labor market. *Handbook of Labor Economics* 3:3143–259.

Athey, S., and G. W. Imbens. 2015. Machine learning methods for estimating heterogeneous causal effects. Working Paper.

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

Bertrand, M. 2011. New perspectives on gender. In *Handbook of Labor Economics*, vol. 4, 1543–90. Elsevier.

Bertrand, M., and E. Duflo. 2017. Field experiments on discrimination. *Handbook of Economic Field Experiments* 1:309–93.

Bertrand, M., E. Kamenica, and J. Pan. 2015. Gender identity and relative income within households. *Quarterly Journal of Economics* 130:571–614.

Bertrand, M., and S. Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review* 94:991–1013.

Blau, F. D., and L. M. Kahn. 2017. The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature* 55:789–865.

Brown, T., B. Mann, N. Ryder, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33:1877–901.

Cassel, J., J. Lerner, and E. Yimfor. 2022. Racial diversity in private capital fundraising. NBER Working Paper.

Chernenko, S., and D. S. Scharfstein. 2022. Racial disparities in the paycheck protection program. Working Paper, Purdue Krannert, HBS.

———. 2023. The limits of algorithmic measures of race in studies of outcome disparities. Working Paper, Purdue Krannert, HBS.

Chernozhukov, V., W. K. Newey, and R. Singh. 2022. Automatic debiased machine learning of causal and structural effects. *Econometrica* 90:967–1027.

Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter. 2020. Race and economic opportunity in the United States: An intergenerational perspective. *Quarterly Journal of Economics* 135:711–83.

Chetty, R., N. Hendren, P. Kline, and E. Saez. 2014. Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics* 129:1553–623.

Chien, C. V., and J. Grennan. 2024. Unpacking the innovator-inventor gap: Evidence from engineers. Working Paper, Berkeley Law, Emory.

Coluccia, D. M., G. Dossi, and S. Ottinger. 2023. Racial discrimination and lost innovation. Working Paper 5.

Cook, L. D., J. Gerson, and J. Kuan. 2022. Closing the innovation gap in pink and black. *Entrepreneurship and Innovation Policy and the Economy* 1:43–66.

Cook, L. D., and C. Kongcharoen. 2010. The idea gap in pink and black. Working Paper, National Bureau of Economic Research.

Cook, L. D., M. Marx, and E. Yimfor. 2022. Funding black high-growth startups. Working Paper, National Bureau of Economic Research.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel. 2003. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics* 118:909–68.

Derenoncourt, E., and C. Montialoux. 2021. Minimum wages and racial inequality. *Quarterly Journal of Economics* 136:169–228.

Enke, B., and F. Zimmermann. 2019. Correlation neglect in belief formation. *Review of Economic Studies* 86:313–32.

Eyster, E., and G. Weizsacker. 2011. Correlation neglect in financial decision-making. Working Paper, LSE Economics, Humboldt University.

Fairlie, R., A. Robb, and D. T. Robinson. 2022. Black and white: Access to capital among minority-owned start-ups. *Management Science* 68:2377–400.

Fairlie, R., and D. T. Robinson. 2023. Racial differences in access to capital for innovative start-ups. *Entrepreneurship and Innovation Policy and the Economy* 2:149–66.

Farre-Mensa, J., D. Hegde, and A. Ljungqvist. 2020. What is a patent worth? Evidence from the US patent "lottery". *Journal of Finance* 75:639–82.

Frakes, M. D., and M. F. Wasserman. 2017. Is the time allocated to review patent applications inducing examiners to grant invalid patents? evidence from microlevel application data. *Review of Economics and Statistics* 99:550–63.

Fryer, Roland G., J., and S. D. Levitt. 2004. The causes and consequences of distinctively black names. *Quarterly Journal of Economics* 119:767–805.

Hochberg, Y., A. Kakhbod, P. Li, and K. Sachdeva. 2023. Inventor gender and patent undercitation: Evidence from causal text estimation. Working Paper, National Bureau of Economic Research.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow. 2019. The allocation of talent and US economic growth. *Econometrica* 87:1439–74.

Hunt, J., J.-P. Garant, H. Herman, and D. J. Munroe. 2012. Why don't women patent? Working Paper, National Bureau of Economic Research.

Jensen, K., B. Kovács, and O. Sorenson. 2018. Gender differences in obtaining and maintaining patent rights. *Nature biotechnology* 36:307–9.

Khetan, V., R. Ramnani, M. Anand, S. Sengupta, and A. E. Fano. 2022. Causal bert: Language models for causality detection between events expressed in text. In *Intelligent Computing: Proceedings of the 2021 Computing Conference, Volume 1*, 965–80. Springer.

Lerner, J. 2002. 150 years of patent protection. *American Economic Review* 92:221–5.

Levy, G., and R. Razin. 2015. Correlation neglect, voting behavior, and information aggregation. *American Economic Review* 105:1634–45.

Li, Z., and S. Liu. 2023. Gender difference in innovation recognition: A textual analysis approach. *Available at SSRN 4527066* .

Misra, D. 2019. Mish: A self regularized non-monotonic activation function. Working Paper, arXiv.

Radford, A., K. Narasimhan, T. Salimans, et al. 2018. Improving language understanding by generative pre-training. Working Paper, OpenAI.

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

Schuster, W. M., M. Marcowitz-Bitton, and D. R. Gerhardt. 2022. The gender gap in academic patenting. *UC Davis L. Rev.* 56:759–.

Shu, T., X. Tian, and X. Zhan. 2022. Patent quality, firm value, and investor underreaction: Evidence from patent examiner busyness. *Journal of Financial Economics* 143:1043–69.

Tabakovic, H., and T. G. Wollmann. 2018. From revolving doors to regulatory capture? evidence from patent examiners.

Veitch, V., D. Sridhar, and D. Blei. 2020. Adapting text embeddings for causal inference. In *Conference on Uncertainty in Artificial Intelligence*, 919–28. PMLR.

Waldfogel, J. 2023. The welfare effect of gender-inclusive intellectual property creation: Evidence from books. NBER Working Paper.

Wang, X., L. Wu, and L. M. Hitt. 2023. Social media alleviates venture capital funding inequality for women and less connected entrepreneurs. *Management Science* .

Yang, S. 2024. Predictive patentomics: Forecasting innovation success and valuation with ChatGPT. Working Paper, Pace Academy.

Appendix A: Proofs

Proof of Proposition 1.

Because all six normal variables are independent, the system $(q, s_1, s_2, s_3)^T$ follows a multivariate normal distribution. Specifically, since all six initial variables are zero mean, q and all three signals are also zero mean. Therefore, the conditional expectation of q on $(s_1, s_2, s_3)^T$ is given by

$$E[q|s_1, s_2, s_3] = \text{Cov}[q, (s_1, s_2, s_3)]\text{Var}[(s_1, s_2, s_3)^T]^{-1} \cdot (s_1, s_2, s_3)^T. \tag{10}$$

Again using the independence of $(q_i, q_u, q_o, \epsilon_q, \epsilon_u, \epsilon_o)$, we solve

$$\text{Cov}[q, (s_1, s_2, s_3)] = (\sigma_i^2 + \sigma_o^2 + \sigma_u^2, \sigma_o^2, \sigma_u^2). \tag{11}$$

Similarly,

$$\text{Var}[(s_1, s_2, s_3)^T] = \begin{bmatrix} \sigma_i^2 + \sigma_o^2 + \sigma_u^2 + \sigma_{\epsilon_q}^2 & \sigma_o^2 & \sigma_u^2 \\ \sigma_o^2 & \sigma_o^2 + \sigma_{\epsilon_o}^2 & 0 \\ \sigma_u^2 & 0 & \sigma_u^2 + \sigma_{\epsilon_u}^2 \end{bmatrix}. \tag{12}$$

Taking the inverse of this and plugging in, we obtain after simplification,

$$\begin{aligned} &\text{Cov}[q, (s_1, s_2, s_3)] \cdot \text{Var}[(s_1, s_2, s_3)^T]^{-1} = \\ &\left(\frac{\sigma_i^2 + \frac{\sigma_o^2 \sigma_{\epsilon_o}^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} + \frac{\sigma_u^2 \sigma_{\epsilon_u}^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2}}{\sigma_i^2 + \frac{\sigma_o^2 \sigma_{\epsilon_o}^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} + \frac{\sigma_u^2 \sigma_{\epsilon_u}^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2} + \sigma_{\epsilon_q}^2}, \right. \\ &\frac{\sigma_o^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} \cdot \frac{\sigma_{\epsilon_q}^2}{\sigma_i^2 + \frac{\sigma_o^2 \sigma_{\epsilon_o}^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} + \frac{\sigma_u^2 \sigma_{\epsilon_u}^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2} + \sigma_{\epsilon_q}^2}, \\ &\left. \frac{\sigma_u^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2} \cdot \frac{\sigma_{\epsilon_q}^2}{\sigma_i^2 + \frac{\sigma_o^2 \sigma_{\epsilon_o}^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} + \frac{\sigma_u^2 \sigma_{\epsilon_u}^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2} + \sigma_{\epsilon_q}^2} \right). \end{aligned}$$

For ease of notation, we write

$$w_1 = \frac{\sigma^2}{\sigma^2 + \sigma^2_{\epsilon_q}}, \quad \sigma^2 = \sigma^2_i + \frac{\sigma^2_o \sigma^2_{\epsilon_o}}{\sigma^2_o + \sigma^2_{\epsilon_o}} + \frac{\sigma^2_u \sigma^2_{\epsilon_u}}{\sigma^2_u + \sigma^2_{\epsilon_u}}$$

such that

$$E[q|s_1, s_2, s_3] = (w_1, \frac{\sigma^2_o}{\sigma^2_o + \sigma^2_{\epsilon_o}} \cdot (1 - w_1), \frac{\sigma^2_u}{\sigma^2_u + \sigma^2_{\epsilon_u}} \cdot (1 - w_1)) \cdot (s_1, s_2, s_3)^T. \tag{13}$$

□

Proof of Proposition 2.

Similarly to the proof of Proposition 1, the conditional expectation of q on $(s'_1, s_2, s_3)^T$ is given by

$$E[q|s_1, s_2, s_3] = \text{Cov}[q, (s'_1, s_2, s_3)]\text{Var}[(s'_1, s_2, s_3)^T]^{-1} \cdot (s'_1, s_2, s_3)^T. \tag{14}$$

The expressions for both matrices are simpler in this case:

$$\text{Cov}[q, (s'_1, s_2, s_3)] = (\sigma^2_i, \sigma^2_o, \sigma^2_u) \tag{15}$$

$$\text{Var}[(s_1, s_2, s_3)^T] = \begin{bmatrix} \sigma^2_i + \sigma^2_{\epsilon_q} & 0 & 0 \\ 0 & \sigma^2_o + \sigma^2_{\epsilon_o} & 0 \\ 0 & 0 & \sigma^2_u + \sigma^2_{\epsilon_u} \end{bmatrix}. \tag{16}$$

In the same manner as before, we solve

$$\text{Cov}[q, (s'_1, s_2, s_3)] \cdot \text{Var}[(s'_1, s_2, s_3)^T]^{-1} = (\frac{\sigma^2_i}{\sigma^2_i + \sigma^2_{\epsilon_q}}, \frac{\sigma^2_o}{\sigma^2_o + \sigma^2_{\epsilon_o}}, \frac{\sigma^2_u}{\sigma^2_u + \sigma^2_{\epsilon_u}}), \tag{17}$$

such that

$$E[q|s'_1, s_2, s_3] = (\frac{\sigma^2_i}{\sigma^2_i + \sigma^2_{\epsilon_q}}, \frac{\sigma^2_o}{\sigma^2_o + \sigma^2_{\epsilon_o}}, \frac{\sigma^2_u}{\sigma^2_u + \sigma^2_{\epsilon_u}}) \cdot (s'_1, s_2, s_3)^T. \tag{18}$$

Note that $w_2' = w_2 \cdot \frac{1}{1-w_1}$. Since $w_1 \in (0,1)$, $w_2' > w_2$, and, similarly, $w_3' > w_3$. Meanwhile, $w_1 = \frac{\sigma^2}{\sigma^2 + \sigma_{\epsilon_q}^2}$ and $w_1' = \frac{\sigma^2 - x^2}{\sigma^2 + \sigma_{\epsilon_q}^2 - x^2}$, where $x^2 = \frac{\sigma_o^2 \sigma_{\epsilon_o}^2}{\sigma_o^2 + \sigma_{\epsilon_o}^2} + \frac{\sigma_u^2 \sigma_{\epsilon_u}^2}{\sigma_u^2 + \sigma_{\epsilon_u}^2} < \sigma^2$, and thus $w_1' < w_1$. □

Table 1. Patent Application Acceptance and Race

This table reports the summary statistics of the acceptance rates of all patent applications filed, sorted by inventor race. Acceptances is a binary variable, equal to 1 for granted applications and 0 for rejected applications. Predicted acceptance rate is the chance of acceptance as calculated by a neural network trained on all input variables. Race is determined by the plurality race calculated based on the Surgeo BIFSG algorithm, geo-coded to Alexandria, Virginia.

Panel A: Average Acceptance Rates by Race.

	Acceptances		
	Mean	SD	N
Plurality Race			
White	0.677	0.468	2,536,203
API	0.707	0.455	718,773
Black	0.649	0.477	78,664
Hispanic	0.606	0.489	37,177
Mixed	0.669	0.471	4,678
Indigenous American	0.613	0.490	75

Panel B: Average Predicted Acceptance Rate by Race.

	Predicted Acceptance Rate							
	Mean	SD	10 Pct.	25 Pct.	Median	75 Pct.	90 Pct.	N
Plurality Race								
White	0.684	0.194	0.400	0.568	0.727	0.836	0.900	2,536,203
API	0.716	0.183	0.451	0.613	0.760	0.857	0.913	718,773
Black	0.668	0.205	0.365	0.537	0.712	0.831	0.897	78,664
Hispanic	0.635	0.219	0.311	0.491	0.675	0.812	0.890	37,177
Mixed	0.698	0.197	0.418	0.589	0.744	0.849	0.912	4,678
Indigenous American	0.672	0.216	0.400	0.522	0.714	0.863	0.912	75

Table 2. Regression of Application Acceptance on Race

This table reports the results of the regressions of application acceptance on inventor race. Inventor race is defined as the average race across inventors as identified by SURGEO. Application Quality is the neural-network predicted chance of application success. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
White	0.028***	0.028***	0.032***	0.014***	0.014***	0.016***
	(16.62)	(11.34)	(17.70)	(8.90)	(7.54)	(9.47)
Black	-0.043***	-0.043***	-0.036***	-0.016***	-0.016***	-0.014***
	(-14.69)	(-7.93)	(-7.43)	(-5.99)	(-4.37)	(-3.66)
API	0.058***	0.058***	0.053***	0.018***	0.018***	0.016***
	(33.11)	(11.53)	(12.54)	(10.69)	(5.79)	(5.42)
Indigenous American	-0.036	-0.036	-0.031	-0.039	-0.039	-0.035
	(-0.67)	(-0.64)	(-0.55)	(-0.77)	(-0.81)	(-0.71)
Mixed	0.020***	0.020***	0.016***	-0.005	-0.005	-0.007
	(2.85)	(3.30)	(2.65)	(-0.81)	(-1.17)	(-1.49)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	3,375,570	3,375,570	3,375,570	3,375,570	3,375,570	3,375,570
Adjusted R^2	0.001	0.001	0.014	0.118	0.118	0.121

Table 3. Public Companies Application Acceptance and Race

This table reports the summary statistics of the acceptance rates of all patent applications filed by public companies, sorted by inventor race. All variables are as defined in Table 1. In Panel C, t-statistics are reported in parentheses, with standard error clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Average Acceptance Rates by Race.

	Acceptances		
	Mean	SD	N
Plurality Race			
White	0.751	0.432	986,560
API	0.754	0.431	403,499
Black	0.755	0.430	29,348
Hispanic	0.749	0.434	10,562
Mixed	0.744	0.436	2,185
Indigenous American	0.600	0.503	20

Panel B: Average Predicted Acceptance Rate by Race.

	Predicted Acceptance Rate							
	Mean	SD	10 Pct.	25 Pct.	Median	75 Pct.	90 Pct.	N
Plurality Race								
White	0.750	0.167	0.507	0.666	0.793	0.875	0.923	986,560
API	0.760	0.161	0.525	0.679	0.802	0.880	0.926	403,499
Black	0.758	0.164	0.519	0.679	0.803	0.879	0.925	29,348
Hispanic	0.750	0.171	0.495	0.662	0.796	0.879	0.928	10,562
Mixed	0.774	0.151	0.566	0.698	0.813	0.887	0.933	2,185
Indigenous American	0.798	0.232	0.366	0.819	0.892	0.918	0.936	20

Panel C: Regression of Application Acceptance on Race.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
White	-0.001	-0.000	0.013***	0.004	0.004	0.008**
	(-0.15)	(-0.03)	(2.82)	(1.58)	(0.95)	(2.12)
Black	-0.008	-0.005	-0.000	-0.001	-0.001	-0.001
	(-0.78)	(-0.56)	(-0.03)	(-0.18)	(-0.08)	(-0.08)
API	-0.001	0.005	0.007	-0.002	-0.001	0.003
	(-0.09)	(0.44)	(1.02)	(-0.77)	(-0.18)	(0.65)
Indigenous American	-0.124	-0.129	-0.200**	-0.167*	-0.168*	-0.191**
	(-1.07)	(-1.14)	(-2.02)	(-1.78)	(-1.90)	(-2.29)
Mixed	-0.005	-0.003	-0.005	-0.020**	-0.020	-0.014
	(-0.36)	(-0.19)	(-0.43)	(-2.05)	(-1.58)	(-1.36)
Firm & Year Fixed Effects	-	-	Y	-	-	Y
Firm Size Control	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	1,301,539	1,301,539	1,301,539	1,301,539	1,301,539	1,301,539
Adjusted R^2	0	0.002	0.053	0.094	0.094	0.111

Table 4. Private Companies Application Acceptance and Race

This table reports the summary statistics of the acceptance rates of all patent applications filed by private companies, sorted by inventor race. All variables are as defined in Table 1. In Panel C, t-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Average Acceptance Rates by Race.

	Acceptances		
	Mean	SD	N
Plurality Race			
White	0.667	0.471	688,473
API	0.667	0.471	145,726
Black	0.647	0.478	20,588
Hispanic	0.640	0.480	9,470
Mixed	0.694	0.461	993
Indigenous American	0.524	0.512	21

Panel B: Average Predicted Acceptance Rate by Race.

	Predicted Acceptance Rate							
	Mean	SD	10 Pct.	25 Pct.	Median	75 Pct.	90 Pct.	N
Plurality Race								
White	0.673	0.189	0.400	0.563	0.712	0.818	0.882	688,473
API	0.673	0.187	0.406	0.560	0.708	0.819	0.885	145,726
Black	0.660	0.191	0.383	0.542	0.696	0.807	0.875	20,588
Hispanic	0.659	0.195	0.379	0.545	0.697	0.810	0.878	9,470
Mixed	0.677	0.187	0.408	0.565	0.718	0.817	0.883	993
Indigenous American	0.595	0.210	0.350	0.432	0.546	0.748	0.868	21

Panel C: Regression of Application Acceptance on Race.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
White	0.020***	0.020***	0.022***	0.010***	0.010***	0.011***
	(5.93)	(6.13)	(6.51)	(3.20)	(3.42)	(3.68)
Black	-0.008	-0.008	-0.003	-0.008	-0.008	-0.006
	(-1.31)	(-0.95)	(-0.38)	(-1.36)	(-1.08)	(-0.89)
API	0.020***	0.020*	0.012	0.010***	0.010	0.007
	(5.65)	(1.89)	(1.23)	(2.96)	(1.35)	(0.96)
Indigenous American	-0.124	-0.124	-0.132	-0.077	-0.077	-0.079
	(-1.20)	(-1.05)	(-1.13)	(-0.78)	(-0.64)	(-0.67)
Mixed	0.046***	0.046**	0.038**	0.033**	0.033**	0.031**
	(3.03)	(2.92)	(2.38)	(2.28)	(2.38)	(2.20)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	865,271	865,271	865,271	865,271	865,271	865,271
Adjusted R^2	0.000	0.000	0.015	0.085	0.085	0.089

Table 5. Individual Inventors Application Acceptance and Race

This table reports the summary statistics of the acceptance rates of all patent applications filed by individual inventors, sorted by inventor race. All variables are as defined in Table 1. In Panel C, t-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Average Acceptance Rates by Race.

	Acceptances		
	Mean	SD	N
Plurality Race			
White	0.600	0.490	861,170
API	0.629	0.483	169,548
Black	0.543	0.498	28,728
Hispanic	0.500	0.500	17,145
Mixed	0.543	0.498	1,500
Indigenous American	0.676	0.475	34

Panel B: Average Predicted Acceptance Rate by Race.

	Predicted Acceptance Rate							
	Mean	SD	10 Pct.	25 Pct.	Median	75 Pct.	90 Pct.	N
Plurality Race								
White	0.619	0.203	0.324	0.484	0.648	0.780	0.861	861,170
API	0.650	0.197	0.364	0.525	0.685	0.806	0.878	169,548
Black	0.581	0.213	0.277	0.434	0.603	0.754	0.846	28,728
Hispanic	0.551	0.223	0.231	0.395	0.570	0.730	0.833	17,145
Mixed	0.601	0.217	0.268	0.478	0.632	0.769	0.859	1,500
Indigenous American	0.645	0.181	0.461	0.563	0.657	0.779	0.863	34

Panel C: Regression of Application Acceptance on Race.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
White	0.058***	0.058***	0.055***	0.026***	0.026***	0.027***
	(19.73)	(14.09)	(14.88)	(9.58)	(8.66)	(10.26)
Black	-0.043***	-0.043***	-0.041***	-0.018***	-0.018***	-0.017***
	(-9.12)	(-6.09)	(-5.98)	(-4.09)	(-3.88)	(-3.43)
API	0.087***	0.087***	0.084***	0.029***	0.029***	0.027***
	(27.84)	(16.00)	(17.12)	(9.76)	(6.70)	(7.19)
Indigenous American	0.134	0.134	0.147	0.080	0.080	0.085
	(1.59)	(1.64)	(1.74)	(1.02)	(1.16)	(1.19)
Mixed	0.001	0.001	0.001	-0.016	-0.016	-0.018
	(0.06)	(0.04)	(0.05)	(-1.34)	(-1.21)	(-1.33)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	1,078,125	1,078,125	1,078,125	1,078,125	1,078,125	1,078,125
Adjusted R^2	0.002	0.002	0.005	0.123	0.123	0.126

Table 6. Patent Application Acceptance and Gender

This table reports the summary statistics of the acceptance rates of all patent applications filed, sorted by majority inventor gender. Acceptances is a binary variable, equal to 1 for granted applications and 0 for rejected applications. Predicted acceptance rate is the chance of acceptance as calculated by a neural network trained on all input variables. Gender is obtained from USPTO data. Patents with equal male and female inventorship are classified as female.

Panel A: Average Acceptance Rates by Gender.

	Acceptances		
	Mean	SD	N
Majority Gender			
Female	0.638	0.481	389,362
Male	0.690	0.463	2,912,399

Panel B: Average Predicted Acceptance Rate by Gender.

	Predicted Acceptance Rate							
	Mean	SD	10 Pct.	25 Pct.	Median	75 Pct.	90 Pct.	N
Majority Gender								
Female	0.652	0.208	0.345	0.514	0.692	0.819	0.892	389,362
Male	0.696	0.190	0.420	0.585	0.739	0.843	0.904	2,912,399

Table 7. Regression of Application Acceptance on Gender

This table reports the results of the regressions of application acceptance on inventor gender. Gender is obtained from USPTO data. Patents with equal male and female inventorship are classified as female. Application Quality is the neural-network predicted chance of application success. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.052***	0.052***	0.055***	0.016***	0.016***	0.018***
	(65.56)	(22.16)	(26.91)	(21.25)	(10.61)	(13.05)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	3,301,761	3,301,761	3,301,761	3,301,761	3,301,761	3,301,761
Adjusted R^2	0.001	0.001	0.015	0.117	0.117	0.120

Table 8. Application Acceptance and Gender by Applicant Type

This table reports the results of the regressions of application acceptance on inventor gender, sorted into three groups: large firms, private companies, and individual inventors. All variables are as defined in above tables. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Large Firms.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.046***	0.044***	0.032***	0.017***	0.017***	0.015***
	(9.96)	(9.20)	(8.90)	(6.31)	(6.19)	(6.89)
Firm & Year Fixed Effects	-	-	Y	-	-	Y
Firm Size Control	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	1,273,966	1,273,966	1,273,966	1,273,966	1,273,966	1,273,966
Adjusted R^2	0.001	0.003	0.046	0.093	0.093	0.109

Panel B: Private Companies.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.046***	0.046***	0.050***	0.016***	0.016***	0.018***
	(28.25)	(17.12)	(18.57)	(10.11)	(7.21)	(8.08)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	844,672	844,672	844,672	844,672	844,672	844,672
Adjusted R^2	0.001	0.001	0.016	0.085	0.085	0.089

Panel C: Individual Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.070***	0.070***	0.069***	0.019***	0.019***	0.019***
	(47.00)	(29.94)	(24.97)	(13.45)	(10.31)	(10.08)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	1,040,739	1,040,739	1,040,739	1,040,739	1,040,739	1,040,739
Adjusted R^2	0.002	0.002	0.006	0.123	0.123	0.125

Table 9. Patent Application Acceptance and Applicant Type

This table reports the summary statistics of the acceptance rates of all patent applications filed, sorted by applicant type (large firm, private firm, or individual inventor). Acceptances is a binary variable, equal to 1 for granted applications and 0 for rejected applications. Predicted acceptance rate is the chance of acceptance as calculated by a neural network trained on all input variables.

Panel A: Average Acceptance Rates by Applicant Type.

	Acceptances		
	Mean	SD	N
Applicant Type			
Public	0.751	0.433	1,465,572
Private	0.667	0.471	846,644
Individual	0.601	0.490	1,075,552

Panel B: Average Predicted Acceptance Rate by Applicant Type.

	Predicted Acceptance Rate							
	Mean	SD	10 Pct.	25 Pct.	Median	75 Pct.	90 Pct.	N
Applicant Type								
Public	0.746	0.152	0.533	0.668	0.782	0.859	0.906	1,465,572
Private	0.655	0.172	0.415	0.553	0.685	0.785	0.850	846,644
Individual	0.618	0.190	0.347	0.491	0.643	0.767	0.848	1,075,552

Table 10. Regression of Application Acceptance on Applicant Type

This table reports the results of the regressions of application acceptance on applicant type (large firm, private firm, or individual inventor). Individual inventor is omitted from the regression due to collinearity, and private/public coefficients may be read as relative to individuals. Application Quality is the neural-network predicted chance of application success. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Large Firms	0.150***	0.150***	0.122***	0.075***	0.075***	0.066***
	(255.92)	(11.15)	(12.56)	(130.94)	(12.50)	(13.33)
Private Firms	0.066***	0.066***	0.033**	0.030***	0.030***	0.021***
	(98.59)	(3.90)	(2.47)	(47.46)	(3.65)	(2.70)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	3,387,768	3,387,768	3,387,768	3,387,768	3,387,768	3,387,768
Adjusted R^2	0.019	0.019	0.026	0.112	0.112	0.115

Table 11. Patent Application Acceptance and Race: Sorted by Gender

This table reports the results of the regressions of application acceptance on inventor race, sorted into two groups by inventor gender. All variables are as defined in above tables. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Female Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
White	0.022***	0.022***	0.024***	0.017***	0.017***	0.019***
	(4.37)	(3.64)	(4.22)	(3.50)	(3.13)	(3.62)
Black	-0.038***	-0.038***	-0.033***	-0.007	-0.007	-0.004
	(-4.75)	(-3.44)	(-3.08)	(-0.93)	(-0.80)	(-0.50)
API	0.113***	0.113***	0.104***	0.037***	0.037***	0.035***
	(22.15)	(10.75)	(11.18)	(7.65)	(4.67)	(4.33)
Indigenous American	-0.144	-0.144	-0.149	-0.047	-0.047	-0.047
	(-1.20)	(-0.87)	(-0.86)	(-0.42)	(-0.29)	(-0.28)
Mixed	0.042**	0.042*	0.035	0.019	0.019	0.019
	(2.02)	(1.93)	(1.57)	(0.96)	(0.93)	(0.94)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	389,362	389,362	389,362	389,362	389,362	389,362
Adjusted R^2	0.009	0.009	0.022	0.131	0.131	0.135

Panel B: Male Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
White	0.024***	0.024***	0.029***	0.012***	0.012***	0.014***
	(13.49)	(8.83)	(13.59)	(7.18)	(5.85)	(7.07)
Black	-0.039***	-0.039***	-0.033***	-0.017***	-0.017***	-0.014***
	(-12.29)	(-7.60)	(-6.92)	(-5.47)	(-4.28)	(-3.63)
API	0.052***	0.052***	0.047***	0.016***	0.016***	0.014***
	(27.34)	(13.10)	(14.99)	(9.12)	(6.91)	(6.50)
Indigenous American	0.005	0.005	0.010	-0.033	-0.033	-0.028
	(0.08)	(0.07)	(0.16)	(-0.57)	(-0.58)	(-0.48)
Mixed	0.013	0.013*	0.010	-0.008	-0.008	-0.009*
	(1.64)	(1.87)	(1.54)	(-1.15)	(-1.58)	(-1.81)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	2,912,399	2,912,399	2,912,399	2,912,399	2,912,399	2,912,399
Adjusted R^2	0.001	0.001	0.014	0.113	0.113	0.117

Table 12. Patent Application Acceptance and Gender: Sorted by Race

This table reports the results of the regressions of application acceptance on inventor gender, sorted into two groups by inventor race. All variables are as defined in above tables. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: White Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.082***	0.082***	0.083***	0.024***	0.024***	0.025***
	(79.40)	(42.78)	(43.70)	(23.94)	(14.20)	(15.07)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	2,499,887	2,499,887	2,499,887	2,499,887	2,499,887	2,499,887
Adjusted R^2	0.003	0.003	0.015	0.116	0.116	0.119

Panel B: Black Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.079***	0.079***	0.077***	0.014**	0.014*	0.014*
	(11.33)	(9.56)	(9.06)	(2.16)	(1.88)	(1.85)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	35,862	35,862	35,862	35,862	35,862	35,862
Adjusted R^2	0.004	0.004	0.015	0.156	0.156	0.160

Panel C: API Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.019***	0.019***	0.020***	0.006***	0.006***	0.007***
	(14.16)	(6.93)	(8.34)	(5.15)	(2.94)	(3.46)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	684,287	684,287	684,287	684,287	684,287	684,287
Adjusted R^2	0.000	0.000	0.014	0.111	0.111	0.114

Panel D: Hispanic Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.080***	0.080***	0.078***	0.026***	0.026***	0.027***
	(15.27)	(11.42)	(12.16)	(5.19)	(3.94)	(4.19)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	77,230	77,230	77,230	77,230	77,230	77,230
Adjusted R^2	0.003	0.003	0.018	0.136	0.136	0.139

Panel E: Indigenous American Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.229*	0.229	0.307	0.002	0.002	0.142
	(1.67)	(1.21)	(1.38)	(0.02)	(0.01)	(0.59)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	73	73	70	73	73	70
Adjusted R^2	0.024	0.024	0.039	0.169	0.169	0.145

Panel F: Mixed Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.051**	0.051**	0.053**	-0.003	-0.003	-0.005
	(2.39)	(2.35)	(2.39)	(-0.17)	(-0.15)	(-0.21)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	4,422	4,422	4,422	4,422	4,422	4,422
Adjusted R^2	0.001	0.001	0.009	0.135	0.135	0.137

Table 13. Discrimination and Gendered Names

This table reports the results of the regressions of application acceptance on inventor gender, sorted into two groups by whether or not the inventor's name is obviously gendered. All variables are as defined in above tables. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Inventors without Gendered Names.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.031***	0.031***	0.033***	0.010***	0.010***	0.012***
	(17.10)	(7.69)	(8.95)	(5.97)	(3.32)	(3.91)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	429,052	429,052	429,052	429,052	429,052	429,052
Adjusted R^2	0.001	0.001	0.014	0.105	0.105	0.108

Panel B: Inventors with Gendered Names.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.077***	0.077***	0.078***	0.023***	0.023***	0.024***
	(57.77)	(36.16)	(36.99)	(18.32)	(12.73)	(13.51)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	1,535,319	1,535,319	1,535,319	1,535,319	1,535,319	1,535,319
Adjusted R^2	0.002	0.002	0.015	0.111	0.111	0.115

Table 14. Inventor Name Frequency and Application Acceptances

This table reports the results of the regressions of application acceptance on name frequency. Name frequency is the natural logarithm of the total occurrences of a name over all patent applications filed plus one. All variables are as defined in the above tables. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Regression of Acceptance on Name Frequency.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Name Frequency	0.006***	0.006***	0.008***	0.004***	0.004***	0.005***
	(38.95)	(5.49)	(11.58)	(24.54)	(6.04)	(9.43)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	3,301,703	3,301,703	3,301,703	3,301,703	3,301,703	3,301,703
Adjusted R^2	0.000	0.000	0.014	0.117	0.117	0.120

Panel B: Regression with Race and Gender Controls.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Name Frequency	0.008***	0.008***	0.009***	0.004***	0.004***	0.005***
	(44.06)	(9.90)	(17.76)	(23.76)	(8.75)	(13.48)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	3,301,703	3,301,703	3,301,703	3,301,703	3,301,703	3,301,703
Adjusted R^2	0.003	0.003	0.017	0.117	0.117	0.120

Table 15. Inventor Location and Application Acceptances

This table reports the results of the regressions of application acceptance on whether inventors live in big cities. A big city is any of the 50 most populous cities in the United States. An application is from a big city if any inventor listed on the application is from a big city. All variables are as defined in the above tables. T-statistics are reported in parentheses, with standard error clustered at the year level where indicated. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Large City	0.019***	0.019***	0.012***	0.008***	0.008**	0.005*
	(28.76)	(6.65)	(5.85)	(13.30)	(2.59)	(2.07)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	2,231,139	2,231,139	2,231,139	2,231,139	2,231,139	2,231,139
Adjusted R^2	0.000	0.000	0.013	0.127	0.127	0.130