Understanding Patenting Disparities via Causal Human+Machine Learning*

Lin William Cong† Stephen Yang‡

First draft: Oct 2024; current draft: Feb 2025

Abstract

We introduce a framework combining advanced AI models with broadly applicable perception tests to overcome restrictive assumptions in causal machine learning and identify high-dimensional discrimination using unstructured data. Applying such causal human+machine learning (C-HML) to patenting and leveraging large-language-model (LLM) embeddings to create objective measures of patent quality, we document systematic bias by the United States Patent & Trademark Office against under-represented groups, especially females and Black innovators, unexplained by the purported criteria for patent approvals and thus likely illegal. Jointly investigating multiple dimensions of bias for the first time, we also discover a novel, large *affiliation bias*—individuals are severely disadvantaged relative to those working under large, public companies—that dominates other disparities. Furthermore, various factors from innovation quality to location can either mitigate or compound disparities, and the interactions of various biases manifest significantly in approval decisions, with misogyny persisting even with advantageous affiliations, for example. We rationalize the findings via a simple model of correlation neglect in which examiners mistakenly interpret traits such as gender, race, and affiliation as providing additional information.

JEL Classification: G30, J15, J16, O31.
Keywords: AI, Causal Inference, Discrimination, Innovation, LLMs, Machine Learning.

*We are grateful to Jillian Grennan, Gerard Hoberg, Josh Lerner, Asaf Manela, and Amit Seru for helpful comments and discussions. We thank Tanvi Kiggar, Rishi Kumar, Jordan Velte, Jonathan Braun, Danish Safdariyan, Vineet Saraf, and Lauren Sutioso for outstanding research assistance.
†Cornell SC Johnson College of Business (Johnson), ABFER, and NBER. Email: will.cong@cornell.edu.
‡Pace Academy. Email: stephen.yang25@paceacademy.org.

Racial and gendered disparities in economic outcomes have been widely documented and constitute pressing concerns around the world (e.g., Bertrand, Kamenica, and Pan, 2015; Blau and Kahn, 2017; Chetty et al., 2014, 2020; Derenoncourt and Montialoux, 2021), given their unfairness, hindrance for growth, and distortion 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 bridge across social groups. However, many already question whether they are, at least historically, largely a false hope due to prevalent discrimination in VC or bank funding (e.g. Chernenko and Scharfstein, 2022; Fairlie, Robb, and Robinson, 2022; Wang, Wu, and Hitt, 2023).

We seek to develop a systematic, robust identification and quantification framework to understand unfair disparities in general, and in the patenting process in particular. Despite previous studies investigating gaps in patenting activity between races and genders, causal evidence is limited and a unifying analysis of the relative effects of different forms of patenting discrimination and their interactions is missing—which we show theoretically to be necessary. Recent applications of some causal machine learning methods show significant progress (e.g., Hochberg et al., 2023), but strong underlying assumptions potentially preclude wider applications. Therefore, we develop a novel causal human+machine learning (C-HML) approach to address these issues. The key idea is that in addition to exploiting the effective non-linear control offered by AI language models, we can use human perceptions of inventor characteristics to identify causal effects relying on an arguably innocuous assumption that perception beyond the actual characteristics is not correlated with any residual innovation quality not captured by the text embeddings. We theoretically demonstrate the potential drawbacks of existing causal ML for identification and how C-HML comes to the rescue.

As for our application of C-HML to patenting, we note that 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 indisputable 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, on aggregate, worse innovation produced by certain groups. Second is a "discrimination" effect, in 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 find that while quality disparities add significantly to the aggregate under-acceptance of minority patents, a majority of the difference between groups seems to be a result of discriminatory treatment.

In the case of disparities in job opportunities and hiring, Bertrand and Mullainathan (2004) conducts a seminal 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.[1] Second, the topic of patent applications, which is critical to the patent acceptance process, has been understudied, largely because of prior data limitations.

We present solutions to both of these problems by leveraging a Generative AI measure of patent application quality built on Yang (2025), which serves an equivalent function to CausalBERT inference (Veitch, Sridhar, and Blei, 2020), a popular causal machine learning methodology, and essentially allows a data-driven non-linear control on quality. Applying LLM technology and machine learning to the patent's contents, we train an objective measure of application quality on historical data of applications and approval outcomes.[2] 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.

This causal ML identification approach requires that correlation between a discriminating dimension (e.g., gender) and innovation quality is fully captured by the AI language model, which

[1] 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 the former is likely illegal.

[2] Instead of using the LLM itself, we apply only the transformer embedding and then train our own secondary models. The embedding is purely a semantic vectorization of the text and mitigates concerns regarding both look-ahead and algorithmic bias.

may not hold. The innovation of C-HML is to add human perception tests on the names of inventors, which significantly relaxes the assumptions needed. 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; for example, male-sounding names experience a significant advantage among female inventors. Similarly, for race, white-sounding names receive preferential treatment (controlling for true race), and Black-sounding names are at a disadvantage. When feasible, applying this technique significantly weakens the otherwise challenging identifying assumption of causal ML because whether or not a name is obviously gendered or racialized is an intrinsic non-economic property of the name. This human perception technique can complement and strengthen the analysis of unfair regulatory bias across many areas involving subjective application review.[3]

Our C-HML approach enables us to 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. We further 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.

Importantly, bias could also manifest in non-demographic characteristics such as firm affiliation, which is new to the literature—examiners may irrationally favor applications filed by Microsoft over individuals. Based on the assignee (legal owner of the patent's intellectual property rights), we label patents as belonging to individual inventors, private firms or public firms. There is a substantial gap between large firms and individuals, with large firm applications experiencing an acceptance rate over 15 percentage points higher. This difference is not driven by skill only; we find significant causal discrimination and preferential bias for public firms, which is even larger than any of the

[3]In reality, assumptions may not hold perfectly for either human perceptions or machine learning approaches, and thus C-HML allows for an over-saturation approach, which is effective except in the unlikely case where several assumptions fail to hold by a wide margin.

demographic dimensions, reaching 6.6 percentage points. Patents are supposedly judged solely on their content, so there is no legal justification for why the same application filed by Apple instead of an individual ought to be accepted instead of rejected.[4] If anything, there exists a compelling argument for preferential treatment towards individual inventors to better encourage innovation by broader subsets of the population.

To the best of our knowledge, despite being highly relevant to discussions of social mobility and entrepreneurship, the affiliation bias is completely new to the literature. We also document several other hitherto unstudied dimensions of discrimination. In particular, names lead to discrimination on the basis of frequency; inventors with rare names are at a substantial disadvantage. Moreover, 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.

Returning to the outsized importance of affiliation bias, we next study cross-sectional variation in demographic discrimination across types of applicants, which is a unique benefit of using the AI-Quality approach rather than previous methodologies. The hypothesis is that 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 could 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. 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.

[4]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.

These observations prompt us to conduct general tests for interaction effects of race and gender too, which is novel to the literature. The findings are 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. Meanwhile, in API inventors, gendered discrimination is substantially reduced: because these names are less obviously gendered to native English speakers, less discrimination occurs.

A simple homophily effect does not explain many of the results. An alternative natural interpretation is that the examiners intentionally discriminate. We, however, do not want to assert this without clear scientific evidence. Instead, we demonstrate in a simple model in Section 6 that discrimination and biases may arise unintentionally as a result of correlation neglect. Because race, gender, and other factors are often 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. While this may not be the only source of discrimination, our model provides a unifying framework that rationalizes all of the dimensions of discrimination and biases observed empirically.

Our first main contribution is to introduce the C-HML framework, which has the potential for a wide range of applications in studying discrimination and bias in general. we are also 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. Equally novel to the innovation literature, we present robust causal evidence that names drive the discrimination effect.

Our paper thus relates to several strands of research. The first concerns economic discrimination against minority groups. 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. Concerning racial disparities in entrepreneurship and business, Aghion et al. (2017) and Bell et al. (2019) study how factors such as neighborhoods "nurturing" could explain the gap in patenting quality. Discrimination against minorities has also been 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). We show yet another avenue of discrimination.

More specifically for disparities in patenting, Tabakovic and Wollmann (2018) and Coluccia, Dossi, and Ottinger (2023) both show potential racial discrimination. Meanwhile, in gendered discrimination, Hunt et al. (2012) attempts to explain why the majority of inventors are male, while Li and Liu (2023) demonstrates an unfair gendered gap in patent citations, Schuster, Marcowitz-Bitton, and Gerhardt (2022) observes a lack of female participation in transfers from academia to industry and Jensen, Kovács, and Sorenson (2018) generally confirms worse proportional outcomes. 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.[5] Moreover, Chien and Grennan (2024) and Chien, Grennan, and Sandvik (2025) investigate voluntary participation in patenting, i.e., the "innovator-inventor gap," and related solutions. We directly study patent review itself.

Our paper also complements an extensive literature studying the unreliability of patent examiners by demonstrating that examiners discriminate based on inventor names. 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 (2025) argues that examiners may prioritize applications with strong writing despite weak underlying technology.

The last point implies that our paper adds to the study of how names relate to economic outcomes.

[5] Lerner (2002) observes that some countries also historically prefer domestic over foreign applicants.

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).[6]

Importantly, we highlight, for the first time, the importance of applicant credibility (via firm affiliation) in patenting, in line with similar results in venture capital. Due to the high risk of entrepreneurship and innovation, decision-makers (venture capital investors) often pay close attention to the identity and credibility of the applicant (Gompers et al., 2020; Bernstein, Korteweg, and Laws, 2017). Particularly, Gompers, Lerner, and Scharfstein (2005) and Hacamo and Kleiner (2022) show that association with large firms can often serve as a strong signal for performance and predict funding outcomes. In venture capital, such decisions are likely justified. In contrast, the legal standards for patenting do not permit consideration of individual characteristics, but examiners nevertheless exhibit bias contributing to large disparities.

More broadly in terms of methodology innovation, this paper contributes to the application of machine learning to studying bias and causal effects, especially involving texts. 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. Causal textual analysis is rare, as revealed in recent comprehensive surveys on the topic (e.g., Hoberg and Manela, 2025). 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 combine powerful LLM embeddings with human perception tests to study high-dimensional causal patenting disparities and their interactions.

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

1. A General Framework for Causally Identifying Discrimination

1.1. Identifying (High-Dimensional) Causal Discrimination

Regulators' decision-making and discrimination. Many important economic outcomes depend on a "regulator"'s evaluation of an application submitted by an individual, ranging from online lending to mortgage applications to patenting.[7] In all of these cases, the quality of the application (perhaps including personal characteristics such as credit score) may be the only legal determinant of outcomes. Quality may correlate with race or other demographic characteristics, and often does. Long-term structural inequities and other factors lead to, for example, aggregate differences in credit score. Nevertheless, a decision-maker should never reject an application that would have been accepted in the counterfactual with a different race or gender *ceteris paribus*.

Clearly, it is difficult but imperative to extract the direct causal effect (if any) of discriminatory factors on outcomes through unfair regulatory/approval treatment. This is particularly troublesome when discriminatory factors correlate with fair standards for the outcome as in the above cases. The econometrician's objective then lies in determining the influence (if any) of said factors on regulators' decisions *beyond* the extent to which they drive actual economic variables. Otherwise, it is impossible to distinguish *bias* from *structural inequity*.

For the ease of exposition, we now describe all variables within the context of patenting, but the findings remain relevant for all applications potentially subject to discrimination. Suppose that the merit, or quality, q of a patent application, i.e., whether it fairly deserves to be approved, is driven by two independent groups of factors such that $q = q_i + q_c$. q_i is determined by the idiosyncratic quality of the application, i.e., the ingenuity and quality of the idea itself. Meanwhile, q_c is driven by the individual characteristics such as race, gender, or socioeconomic background.[8] Akcigit, Grigsby, and Nicholas (2017), Bell et al. (2019), and others document the importance of environmental factors in nurturing innovative skills. Also potentially included in q_c is current employment status: an inventor affiliated with Apple and with access to a successful patent lawyer will, on average, produce

[7] We generically refer to the decision-makers who may discriminate as "regulators."

[8] In the baseline model, we simplify individual characteristics to this single variable q_c. In the following section, we consider the case of multiple, interacting factors such as race and gender.

higher-quality inventions.

The patent examiner's legal obligation is to make decisions based on their judgement of q; essentially, they must attempt to estimate q by reading the application. In reading the application, the examiner obtains a noisy signal $s_1 = q + \epsilon_1$, with ϵ_1 independent from q, q_i, and q_c. Note that because examiners are time-constrained and must make subjective judgements, ϵ_1 can potentially be significant. In addition, examiners can observe, through the application, $s_2 = q_c + \epsilon_2$. If q_c denotes gender, for example, examiners can infer gender from the listed name of the inventor.

Legally, examiners must make a decision based only on s_1, i.e., their reading of the application. Modeling examiner decisions as $y = \beta_1 s_1 + \beta_2 s_2$, we can draw the following correspondences:
$\beta_2 = 0$ **(no discrimination):** The examiner indeed follows legal requirements and makes the same decision as if they had not observed the individuals gender, race, and so on.
$\beta_2 \neq 0$ **(discrimination):** A biased examiner illegally discriminates, incorporating the characteristics of the individual into the decision.

Now, consider the role of the empiricist studying the examination process, who seeks to determine whether unfair discrimination is taking place. Optimally, one could conduct a large-scale randomized controlled trial, holding s_1 constant and varying s_2 to observe outcomes for identical applications filed by, for example, inventors of differing genders. Nevertheless, most settings lack appropriate conditions for such experimental techniques. Thus, the econometrician must investigate the regulator's (lack of) biases using only ex-post observation of the examiner's ultimate decisions Y. In other words, the econometrician must estimate β_2 without knowledge of s_1 and s_2, which are subjective signals the examiner determines based on priors, experience, and so on.

Causal ML for estimating β_2? Often, an econometrician has access to data regarding q_c. In this section, we take the example where q_c is gender, for ease of exposition. An econometrician may, in attempting to determine the presence of gendered discrimination, run the following regression:

$$y = \hat{\beta}_2 q_c + \epsilon. \tag{1}$$

Nevertheless, estimating such a model leads to biased results.

Proposition 1. *Even if* $\beta_2 = 0$*,* $corr(y, q_c) \neq 0$ *and* $\hat{\beta}_2 \neq 0$*.*

The proofs of this and later propositions are provided in the Appendix. Essentially, regardless of true discrimination, the econometrician always observes an association between outcomes and gender, because fair quality q is related to q_c. Clearly, controlling for q would allow the econometrician to obtain a better estimate $\hat{\beta}_2$, but q is also unobservable.

Machine learning techniques present a potential solution to this identification issue. In patenting (as with most settings), the entire text of the application is available to the econometrician. By using a textual machine learning model to analyze the application, the econometrician can generate a noisy signal $q_{ML} = q + \epsilon_{ML}$, where the variance of ϵ_{ML} decreases with the quality of the model. If ϵ_{ML} is identically zero, then the machine learning model is perfect. Otherwise, let ϵ_{ML} be distributed with zero mean and variance σ^2_{ML}. With this signal, the econometrician can run the following model:

$$y = \hat{\beta}_1 q_{ML} + \hat{\beta}_2 q_c + \epsilon. \tag{2}$$

Proposition 2. *If* $corr(\epsilon_{ML}, q_c) = 0$*, the estimation error is given by the following expression:*

$$\hat{\beta}_2 - \beta_2 = \beta_1 \frac{\sigma^2_{ML}}{\sigma^2_{q_i} + \sigma^2_{ML}}.$$

If the machine learning model is highly accurate and unbiased (i.e., $\sigma^2_{ML} \approx 0$), then it essentially "controls" for quality, allowing the econometrician to identify any illegal consideration of q_c and extract the value of β_2. However, σ^2_{ML} may not always be sufficiently small for accurate estimation of β_2. This will lead to an over-estimation of the effect and could even cause the identification of false positives even if $\beta_2 = 0$. Additionally, $corr(\epsilon_{ML}, q_c) = 0$ may fail to hold in some scenarios. If $corr(\epsilon_{ML}, q_c) \neq 0$, it is possible to not only over-estimate β_2, but in fact flip the sign, making it difficult to draw conclusive results.

The C-HML solution to estimating β_2. Under the above formulation of the econometrician's problem, it is indeed very difficult (absent perfect machine learning) to always extract an unbiased $\hat{\beta}_2$. However, let us return to the assumption above that the econometrician observes only q_c and not s_2.

To be clearer in the case of inventor gender in patenting, q_c is the true gender of the applicant, while s_2 is the examiner's opinion of the inventor's gender (since true demographics are not available). It seems that the subjective s_2 should be unobservable, but, in fact, it is possible for the econometrician to approximate this belief ex-post.

Recall that the examiner uses s_2 because q_c is not available, and the examiner is inferring characteristics from the application. Specifically, if examiners consider gender in their decisions, they must do so based on inventor names. Thus, to be more precise, s_2 is actually $s_N = q_c + \epsilon_N$. By observing the names (which are also available ex-post), the econometrician can individually and carefully judge the impressions the names make on a reader, i.e., s_N. For gendered discrimination, this essentially means the "male-ness" of a name.

Now, a critical point to notice is that ϵ_N, the difference in "perceived gender" from actual gender, ought to be non-economic in nature, and thus independent of other variables. While many discriminatory factors such as race or gender are economic, the appearance of a name is not. That is, whether one has a gender-neutral name should not correspond to any fundamental differences across applications, *holding gender constant*. With this in mind, the econometrician can incorporate s_N, the *human perception* variable into the regression model:

$$y = \hat{\beta}_1 q_{ML} + \hat{\beta}_2 s_N + \epsilon, \text{holding } q_c \text{ constant.} \tag{3}$$

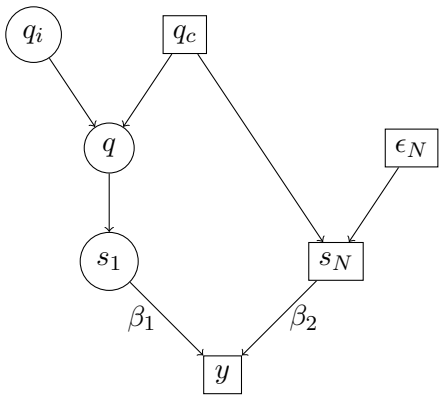
Proposition 3. *If* $corr(\epsilon_N, q_{ML}) = 0$*, then* $\hat{\beta}_2 = \beta_2$.

Essentially, if the examiner is not discriminating and the correlation between y and q_c occurs only because of q_c's relation to q, then ϵ_N should have no effect on Y. The key is that examiners only observe q_c indirectly through the non-economic variable of names, which allows for a clean identification leveraging only the variation in ϵ_N.

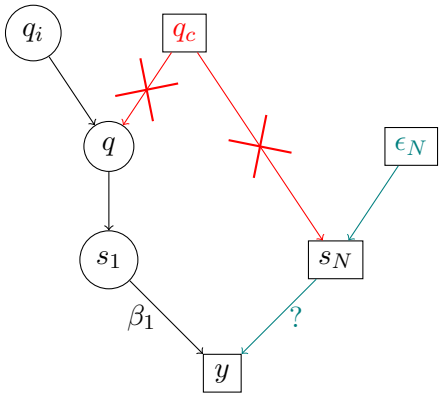
Figure 1 visualizes these relations between variables in the model. Squares represent observable features, while circles are latent. Again, as shown in Panel B, the econometrician eliminates any residual relation between s_1 and s_N by holding q_c constant.

In the following discussion, we consider an extension of the model with multiple discriminatory

Figure 1. Model Illustration



Panel A: Variable Relations



Panel B: ϵ_N Conditional on q_c

factors (i.e., race and gender).

High-dimensional discrimination and mis-measurements. In reality, more than one discriminatory factor may affect q_c. For example, race and gender could both simultaneously play into examiner decisions. Thus, we extend the model to enrich q_c. Specifically, let

$$q_c = \beta_G\, G + \beta_R\, R + \beta_{GR}\, (G \times R), \tag{4}$$

where G is a binary indicator for gender (e.g., 1 if female, 0 if male) and R is a binary indicator for race (e.g., 1 if non-white, 0 if white). Such a setup captures the separate effects of race and gender as well as intersectionality, which can potentially be of outsized importance.

Then, y may be written as

$$y = \beta_1 s_1 + \beta_2 s_N = \beta_1 q_i + \beta_2 \left(\beta_G\, G + \beta_R\, R + \beta_{GR}\, (G \times R)\right) + \beta_1 \epsilon_1 + \beta_2 \epsilon_N \tag{5}$$

$$= \beta_1 q_i + \gamma_G G + \gamma_R R + \gamma_{GR}(G \times R) + \beta_1 \epsilon_1 + \beta_2 \epsilon_N \tag{6}$$

Now, a significant risk arises. Assume that a hypothetical econometrician has a completely robust identification strategy in access to a flawless machine learning model such that $q = q_{ML}$. Even in this perfect situation, the econometrician may still misidentify the presence or magnitude of bias: suppose now that this econometrician focuses only on the gender component of discrimination and

estimates:

$$y = \hat{\beta}_1 q_{ML} + \hat{\gamma}_G G + \eta, \tag{7}$$

If q_c depended only on gender G, then the econometrician would extract the correct value of γ_G without issue as per Proposition 2. Nevertheless, the econometrician's approach is flawed because it ignores R and the interaction $G \times R$.

Proposition 4. *If the econometrician estimates a model considering only G, the estimated coefficient $\hat{\gamma}_G$ will be subject to bias such that:*

$$\hat{\gamma}_G - \gamma_G = (\gamma_R + \gamma_{GR}) \frac{\text{Cov}(G, R)}{\text{Var}(G)}. \tag{8}$$

If $\text{Cov}(G, R) \neq 0$, the estimated effect may be significantly biased and mis-estimated due to omitted variables.[9] The magnitude of the mis-estimation depends on the other discriminatory effects and the extent of the correlation, and thus may easily be large enough to *flip the sign*. Critically, without considering the other factors, there is no way for the econometrician to identify this bias or estimate its potential effects, *even given an otherwise perfect identification strategy*. The econometrician cannot accurately draw any conclusion whatsoever regarding the potential presence of gendered discrimination. In sum, if applications are affected by a combination of discriminatory factors, researchers ought to consider the high-dimensionality of discrimination, accounting for various contributing factors and intersectionality. Our C-HML approach does exactly that.

1.2. The Causal Human+Machine Learning Approach

We now describe the details of C-HML, starting with its machine learning part. We also discuss how C-HML relates to and complements other methodologies.

ML via a AI-based quality measure. Recall that the ultimate goal empirically lies in distinguishing racial disparities driven by structural inequities from unfair discrimination on the part

[9] Why should gender and race be correlated? Obviously, in the population, this is not necessarily the case. But inventors are not a random sample of the population and equilibrium entry is affected by a variety of factors. Both gender and race play a role in this, e.g., white men may have the easiest path to innovation in the United States.

of the regulator. This is because while both present significant social challenges, they necessitate significantly different solutions.

For patent applications, all legally relevant 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. The identifying assumption here is that the measure must accurately determine quality and that the text contains all relevant information, as in Proposition 2.[10]

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 (2025). 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 (Radford 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, 2025). 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 an 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):

$$Mish(x) = x \cdot \tanh(\ln(1 + e^x)). \tag{9}$$

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

[10]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.

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.

Note that LLM approaches can be subject to look-ahead bias (e.g., He et al., 2025. The approach in this papers differs substantially from direct querying of ChatGPT, as employed by Jha et al. (2023), Kim, Muhn, and Nikolaev (2023), Hansen and Kazinnik (2023), and others. Instead, as in Yang (2025), we use only the textual embedding model and then train secondary models on the embedding. This general semantic encoding is unlikely to contain granular information related to specific patenting outcomes. Prompting the model can lead to information retrieval from the dataset, but the embedding simply captures the meaning of the text. Moreover, look-ahead biases would likely skew results towards 0: for example, if the model knows all ex-post outcomes, the control will subsume all variation. In all, it is unlikely that biases in the model could cause incorrect identification of discrimination.

Human perception tests. Next, we describe the human element in C-HML. 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 as shown in Section 1.1. 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.

As with q_c in the model, actual demographic information is not available to patent examiners at

the time of review precisely because of discrimination concerns. Racial and gendered discrimination then happens, if at all, through the examiners' perception of names, i.e., s_N. 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, perceived gender should make no difference, and only the true gender should be related to outcomes. Thus, we can identify whether the examiner is 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, as shown in Proposition 3. In the case of gender, we separate the sample into true male and female inventors, then regress outcomes within each sample on perceived gender (still applying the causal ML control). 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.

Relating to existing methodologies. Finally, we relate C-HML to existing methodologies, including some causal ML approaches. The ideal way to determine the presence of discrimination is performing a 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 cleanly determined that, for example, hiring managers are changing their 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, known as Causal-BERT. 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 the ML Quality approach in our model and that we apply, namely that the text contains all relevant information and that the ML model captures all textual information.

For example, suppose that there is no unfair discrimination and the observed disparities are a result of a confounding factor (such as social upbringing) driving a relation between gender and patent acceptances. Take the illustrative case where the ML model captures zero idiosyncratic information from the text, and thus always "predicts" the mean. Then, under both the Causal-BERT and AI-Quality methodologies, the results will falsely indicate discrimination: in Causal-BERT because the sample means are different across genders, and in AI-Quality because the measure doesn't control for quality variation. Meanwhile, when the ML is perfect, both methods prove causality. In sum, the two techniques require precisely the same assumptions for causal identification. The first difference in the two ML-based approaches lies in their uses, wherein AI-Quality allows for the analysis of high-dimensional factors, interactions, and continuous independent variables rather than only singular binary treatments. As shown in Proposition 4, such considerations can be crucial.

Specifically, Causal-BERT and similar methodologies also face several additional challenges resolved by our unconditional quality approach. In particular, training counterfactual models becomes prohibitively time-consuming 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.[11] Whereas earlier techniques are usually restricted to determining a binary treatment effect, our quality control applies to a wide range and variety of

[11] 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.

disparities. Further, our methodology also allows for determining continuous effects, such as the name frequency results in Section 4.2.

Broadly, the C-HML approach is key because no causal identifying assumption can ever hold perfectly. For example, while Causal-BERT models may have high accuracy at some times, any imperfection could still admit omitted variable bias. Critically, the assumptions for human perceptions are very likely to hold, whereas sophisticated ML models are also generally very accurate. Combining both of these thus significantly raises the bar for identifying a spurious correlation, since they mutually mitigate bias concerns.

2. Institutional Background and Data

2.1. Patenting

When an inventor or firm discovers a new invention, they must file an application with the US Patent and Trademark Office (USPTO).[12] An application consists of a detailed description of the invention, an abstract, specific technical claims, and sometimes figures. Additionally, the examiner also receives a limited set of more structured information: the number of claims, the affiliation of the inventor, and the applicant's name.

Once the application is received, it is randomly assigned to an examiner within the art unit (subject area) of the patent. Then, the examiner should determine whether the application meets the three legal standards of acceptance: originality, usefulness, and non-triviality.[13] Importantly, all of these standards pertain only to the content of the application. Thus, the examiner cannot legally consider external factors such as the identity or affiliation of the applicant in making a decision.

2.2. Data Description

Our primary data source is the USPTO's PatentsView service, which provides bulk textual data for all utility patents granted after 1976 and all patent applications published after 2001. Firm

[12]Patents fall under three categories: utility, design, and plant. This paper focuses on utility patents, which are the most important and by far most common subset.

[13]Patent Act of 1952, 35 U.S.C. § 101-103, and MPEP 2107. A technically fourth statutory (subject matter eligibility) requirement is: the patent must be a process, machine, article of manufacture, or composition of matter.

market capitalization is obtained from CRSP and firm variables from Compustat.

2.3. Variables

AI Quality is produced as described above in Section 1.2. Our measure of individual inventor race is constructed using the Bayesian Improved First Name Surname Geocoding (BIFSG) algorithm via the Surgeo Python package.

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, each individual is first assigned the race to which they most likely belong. After this, if the application has multiple inventors, the plurality is taken again. Gender is taken directly from the USPTO.

Human perceptions, i.e., "Male-ness," "White-ness," and "Black-ness" of names are determined manually by reading names and identifying their impression. For gender, only first names are used, while first name-last name pairs are identified for race.

3. Racial and Gendered Disparities

3.1. Suggestive Evidence of the Racial Gap

One motivation for this study lies in the significant disparities found across races throughout the innovation chain. 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 ex-ante 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 are, on average, 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. 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]

3.2. Identifying Unfair Outcomes

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):[14]

$$Accepted_{i,t} = \sum_{R \in \{Races\}} \beta_n \cdot R_{i,t} + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{10}$$

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

[14] 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, when we say that being White grants an unfair advantage, this is relative to Hispanic inventors.

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, 2025, 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.

3.3. 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 3 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%.[15]

[Include Table 3 Here]

[15]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."

Next, we perform the same quality-controlled test as in Section 3.2. Specifically, we estimate the following regression:

$$Accepted_{i,t} = \beta \cdot Male_{i,t} + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}, \tag{11}$$

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

Table 4 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. Economically, 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 4 Here]

3.4. Confirming Causality via Human Perception

The above results demonstrate the persistence of substantial disparities between male and female and races. Nevertheless, as discussed in Section 1, the AI techniques are potentially subject to bias. In this section, we seek to show that examiners are directly inferring (and discriminating based on) gender and race from *inventor names*. As demonstrated in the model, we exploit the difference between perceived gender and race (from names) and their true values.

Gendered discrimination. First, we seek to rigorously show causality for the gendered discrimination finding. As our "true" value, we use the USPTO provided data, which uses ML to accurately infer inventor gender.[16] Meanwhile, we also manually label the $15,000$ most common inventor names as obviously male, obviously female, or not obviously gendered. By labeling inventors using the impressions their names make on average humans, we extract the examiner's beliefs about gender.

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

In general, human perceived gender often agrees with true gender, but the Pearson correlation coefficient is only around 71%. While the name-based perception is far from random, it also often disagrees with the true value, providing the statistical power needed for our tests.

Essentially, we seek to isolate and remove the underlying correlation of gender and application quality, instead determining if the examiner is biased. In this case, we first split the sample into two: male inventors, and female inventors. Then, within each sample, we perform a similar gendered discrimination regression to Section 3.3, using *perceived* instead of actual gender, and taking only names that are identified as strongly gendered. If the examiner is using only the legal standards for acceptance, then perceived gender, conditional on actual gender, should have no effect on outcomes.

Table 5 reports the regression results. In both samples, the "perceived male-ness" of a name has a significant effect on acceptances, even after controlling for quality. Among female inventors, those with male-sounding names experience an increase of 2 percentage points in acceptance rate, and a similar effect is seen among male inventors. The significance of these coefficients is strong evidence that examiners are not simply examining other aspects of the application correlated with gender but instead are making judgments based on an impression of gender derived from the inventor's name.

[Insert Table 5 Here]

Note that many individuals are ambiguously gendered to the average reader, and are dropped from the above test to isolate the effect of strongly gendered names. These failures to label gender at all result both from gender-neutral names such as "Jordan" but also from cultural differences: American readers are unable to identify names such as Sung-Min as either male or female. Thus, we also perform a test by splitting the sample into the following two subsamples. First, inventors with names obviously gendered to the average American English-speaking reader. Second, inventors without obviously gendered names.

Then, we again perform a gendered discrimination regression. Table 6 reports the results. Although males experience an advantage in both samples, the effect is universally twice as large in inventors with obviously gendered names.[17] This difference is further evidence that examiners tend

[17]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

to reject applications with female-sounding names *ceteris paribus*.

[Insert Table 6 Here]

Racial discrimination. Next, we apply similar techniques for the case of race. Racial discrimination requires a more complicated treatment because there are many potential races, and the human perception error is unclear and may often be correlated with true race. For example, it is much more likely to falsely label a white name as Black than an API name as white. Thus, we perform two tests. First, to identify if white inventors experience an advantage (relative to others). Second, to determine what disadvantage (if any) Black inventors face. For white inventors, we similarly split the sample into white and non-white inventors, then regress acceptances on perceived white-ness.

Table 7 shows the results. Even among white inventors, name white-ness has a significant effect on success. The effect is ven more pronounced among non-white inventors, with a white-sounding name leading to a 2.2 percentage point increase in acceptance rate *controlling for quality*. As with gender, the perceived white-ness of a name should have no underlying link with economic variables. Thus, the examiner must also infer race from names.

[Insert Table 7 Here]

Similarly, we perform the same tests for perceived Black-ness. Table 8 reports these regressions. Among non-Black inventors, there is a small effect that vanishes when controlling for quality. However, among Black inventors, the perceived Black-ness of a name has a very significant negative association with outcomes—Black inventors with more Black-sounding names are more than 3% less likely to receive acceptances. This effect is measured after controlling for race and application quality and thus shows that examiners discriminate substantially against inventors with "Black names."

[Insert Table 8 Here]

In all, the evidence suggests that examiners (whether consciously or unconsciously) make biased directions based on name-based information. Thus, switching to a double-blind procedure could thus significantly cut unfair discrimination.

cause for the persistence of the effect in ambiguously gendered names. Regardless, our result effectively provides a lower bound on the actual impact.

4. Novel Discrimination Factors

4.1. Applicant Types and Affiliation Bias

Beyond race and gender, could other factors also lead to unfair outcomes? Examiners could discriminate (intentionally or unintentionally) based on any relevant individual characteristics. Thus, bias across non-traditional dimensions could also be economically significant. In fact, the degree of bias could increase with the relative importance of the factor in the first place—meaning that certain characteristics could generate even more extreme disparities. One of the most important non-demographic traits of an inventor is *firm affiliation*; essentially, whether or not the patent is associated with a company and what that company is.

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.[18] This can potentially be incredibly harmful in terms of social mobility since poorer entrepreneurs and inventors will obviously file applications as individuals.

To test for the presence of such effects, we define three types of patent applicants: public firms, private firms, and individual inventors. 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 own name, that will belong entirely to themselves if accepted.

Table 9 reports the aggregate acceptance rates and mean application quality across these 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

[18]Whether such a bias is efficient is not immediately clear, but it seems rather inappropriate *and illegal* 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.

individuals, that would be cause for significant concern.

[Insert Table 9 Here]

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

$$Accepted_{i,t} = \beta_1 \cdot Public_{i,t} + \beta_2 \cdot Private_{i,t} + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{12}$$

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 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.[19] A strictly homophily based view of the world may 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. In fact, this *affiliation bias* is of larger magnitude than either racial or gendered discrimination.

4.2. Other Dimensions of Bias

Broadening our examination of discrimination beyond demographic groups and homophily effects also opens the door to several other dimensions. Beyond affiliation with a powerful institution, as discussed above, many other "soft" variables are also available to the examiner. For example, we have

[19] 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.

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:

$$Accepted_{i,t} = \beta \cdot Name\ Frequency_{i,t} + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{13}$$

Panel A of Table 11 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 11 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 11 Here]

Another potential source of bias lies in an inventor's location, which is also reported on the patent application. 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:

$$Accepted_{i,t} = \beta \cdot Large\ Cityi, t + \gamma \cdot Controls_{i,t} + \delta_t + \epsilon_{i,t}. \tag{14}$$

Table 12 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, 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 12 Here]

5. Heterogeneity and Intersectionality in Patenting Disparities

5.1. Cross-Sectional Disparities

5.1.1 Racial Discrimination

In Section 3, we demonstrate the persistence of causally unfair outcomes. However, such examiner biases may not be homogenous. Many soft, unregulated factors including, but not limited to, race and gender, may influence examiner decisions. For example, firm affiliation plays an even larger role in driving biased outcomes. Psychologically, examiners may believe that if the patent is filed by a large firm such as Apple or Microsoft, it likely contains novel, useful technology. An examiner could see this firm affiliation before reading the text of the patent or the name of the inventor, thus "overriding" demographic biases. [20]

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. That is, affiliation with a big name firm may be powerful enough to push the application over the line alone, before the examiner's racial biases even have a chance to come into play. To test this hypothesis, we again split the sample into the three cross-sections: public firms, private firms, and individuals.

Public firms. In Table 13, 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

[20] Or, examiners could prioritize reading the name of the inventor first.

acceptance rate and application quality.[21] In the regression setting, white inventors still experience some advantages, but the effects are much smaller and only appear in a few specifications.[22] Clearly, the discrimination effect is much less prominent for applications filed by public firms.

[Insert Table 13 Here]

Private firms. Table 14 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 14 Here]

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. While an inventor at Apple is unlikely to lose their job simply because of a 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 15 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

[21] 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.

[22] 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.

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 15 Here]

In general, such cross-sectional variation is critical to acknowledge. In this case, failing to filter by affiliation bias would lead one to under-estimate the discriminatory effect. However, if, for example, for inventors affiliated with public firms, the discriminatory effect were reversed (i.e., preference for Black inventors), then a study of the entire sample would reveal no discrimination even though individuals face unfair outcomes. Thus, considering multiple dimensions simultaneously is critical.

5.1.2 Gendered Discrimination

The natural question that then arises is whether gendered discrimination follows the same patterns. 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 quality-controlled regression of application acceptance and gender for all three applicant types (Large Firms, Private Firms, and Individual Inventors).

Table 16 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 difference is negligible. Essentially, female inventors experience a substantial disadvantage *irrespective* of their affiliation with a large firm, 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).

[Include Table 16]

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. For example, gendered bias could be "stronger," not in terms of magnitude but rather persistence.

5.2. 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, as done above. 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 17 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 17 Here]

Table 18 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%). Recall that a large portion of discrimination is driven entirely by name-inferred gender. Since API names are often harder to accurately identify as male or female for the average American, the effect is less present.

[Insert Table 18 Here]

5.3. Quality-Heterogeneous 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.

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 2 to 4.

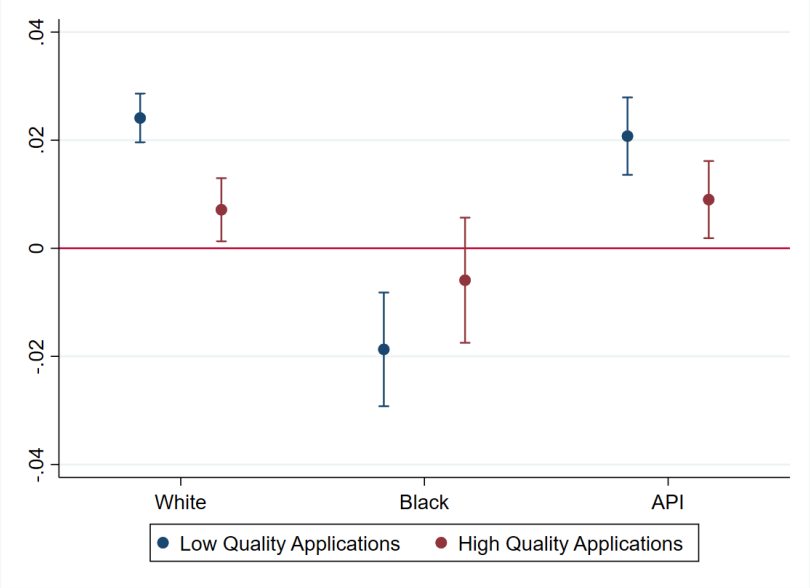


Figure 2. Race Disparities: High- vs Low-Quality Applications

Figure 2 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

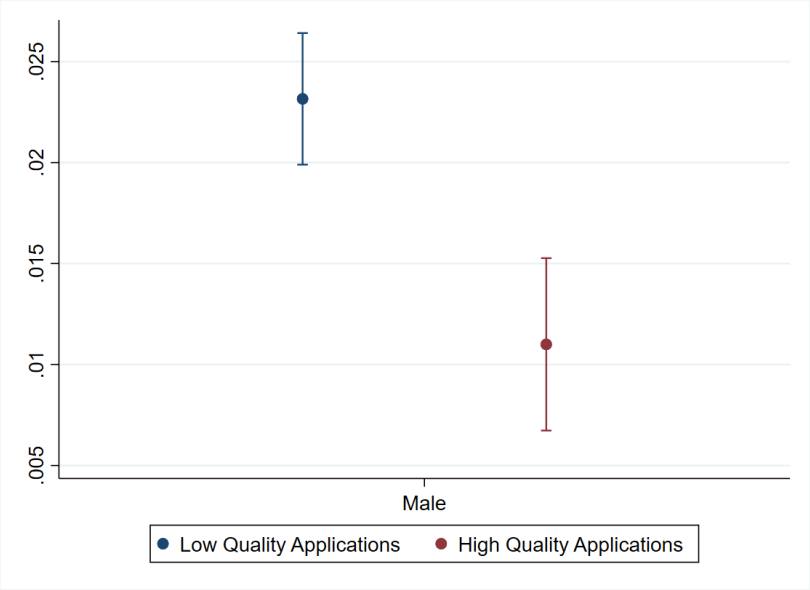


Figure 3. Gender Disparities: High- vs Low-Quality Applications

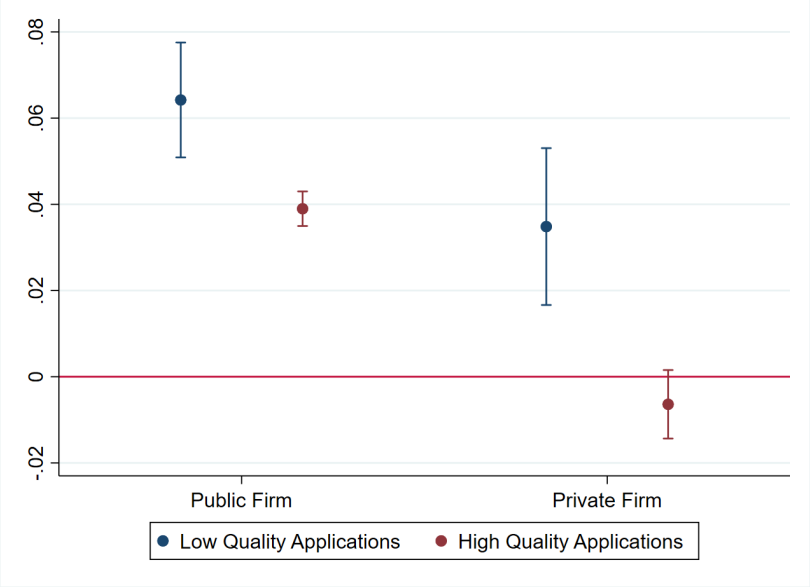


Figure 4. Disparity in Applicant Types: High- vs Low-Quality Applications

insignificant from zero.[23]

Similarly, Figure 3 reveals that the gender disparity is also significantly greater for low-quality patent applications. Figure 4 finds that large firms enjoy substantial advantages for all applications

[23]The difference in the proportional effect is, in fact, greater because low-quality applications have a lower predicted acceptance rate.

filed, with a greater advantage for low-quality applications (relative to individual inventors). Private firms only experience an advantage for low-quality applications.

6. A Theory of Discrimination Due to Correlation Neglect

In Section 1.1, we considered ways in which econometricians can ex-post identify unfair discrimination, which we then applied throughout the paper, indeed confirming the presence of causally unfair disparities. Nevertheless, from the ex-ante perspective, why do decision-makers such as patent examiners exhibit these unfair biases in the first place? Can we explain such decisions in the context of our theoretical framework? Such discrimination is surprising because it is not only inefficient but also poses a career risk to the examiners if discovered. One possible explanation is that the regulators hold intrinsic personal biases against certain groups and intentionally discriminate, such that $\beta_2 > 0$. Such a claim is speculative and too strong to make without supporting evidence. To rationalize the empirical regularities, we propose an alternative explanation in which *unfair* disparities could arise even unintentionally. Our simple model of patent examiner's decision-making involves the concept of correlation neglect. Put simply, examiners may overweight factors such as affiliation, race, or gender because they mistakenly interpret them as independent signals of application strength, leading to "accidental" discrimination.

Recall from Section 1.1 that $q = q_i + q_c$ and that the examiner's job is to estimate the unobservable q. The examiner can read the content of the application, which provides a noisy signal $s_1 = q + \epsilon_1$. If the examiner is perfectly skilled and has full information, $\epsilon_1 = 0$. In practice, this is never the case, and thus, the examiner must seek other signals, namely, $s_N = q_c + \epsilon_N$ in the model. In the earlier model, the examiner's problem of optimally approximating q is considered exogenously solved, and the solution $\beta_1 s_1 + \beta_2 s_N$ is taken as given. Now, we return to show the details of how an examiner might determine the weights β_1 and β_2. For tractability, we add the additional restriction that these independent variables $(q_i, q_c, \epsilon_1, \epsilon_N)$ are normally distributed with zero mean and variance given by $(\sigma_i^2, \sigma_c^2, \sigma_{\epsilon_1}^2, \sigma_{\epsilon_N}^2)$.[24] The examiner then determines $E[q|s_1, s_N]$.

[24]This assumption is not critical for the intuition of the result but is necessary in some form to write interpretable, explicit expressions for the upcoming propositions.

Proposition 5. *Given s_1 and s_N, the conditional expectation of patent quality q is:*

$$E[q|s_1, s_N] = \beta_1 s_1 + \beta_2 s_N$$
$$\textit{where} \qquad \beta_1 = \frac{\sigma^2}{\sigma^2 + \sigma^2_{\epsilon_1}}, \quad \sigma^2 = \sigma^2_i + \frac{\sigma^2_c \sigma^2_{\epsilon_N}}{\sigma^2_c + \sigma^2_{\epsilon_N}},$$
$$\beta_2 = \frac{\sigma^2_c}{\sigma^2_c + \sigma^2_{\epsilon_N}} \cdot (1 - \beta_1).$$

Importantly, if σ_{ϵ_1} is small relative to σ^2_i, then β_1 will approach 1 and β_2 will approach 0. Essentially, if the examiner is good and the innate quality of the idea is the most important, then the examiner should almost use only the text of the application in their estimation of q, i.e., $E[q|s_1, s_N] \approx E[q|s_1]$. If this were the case, β_2 would likely be too small to detect empirically.

One reason this may not hold in reality is that the examiner may easily overlook the correlation between the signals s_1 and s_N. This is a well-documented effect known as *correlation neglect*, where agents do not understand that distinct signals contain repetitive information. DeMarzo, Vayanos, and Zwiebel (2003), Eyster and Weizsacker (2011), and Enke and Zimmermann (2019) study the causes and effects of similar biases in various settings.[25] 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 signals as $s'_1 = q_i + \epsilon_1$ and $s_N = q_c + \epsilon_N$. Essentially, the examiner thinks that they are gleaning additional information about quality from the affiliation, name, and other characteristics of the inventor even though it is already contained in the actual content of the invention.

Proposition 6. *1) When the examiner misinterprets s_1 as s'_1 under correlation neglect assumptions, the conditional expectation of patent quality q is given by:*

$$E[q|s'_1, s_N] = \beta'_1 s'_1 + \beta'_2 s_N$$
$$\textit{where} \qquad \beta'_1 = \frac{\sigma^2_i}{\sigma^2_i + \sigma^2_{\epsilon_1}}, \quad \beta'_2 = \frac{\sigma^2_c}{\sigma^2_c + \sigma^2_{\epsilon_N}}.$$

2) The following always holds: $\beta'_1 < \beta_1, \quad \beta'_2 > \beta_2$.

[25] 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.

The correlation neglect of the examiner results in an overestimation of the importance of s_N, 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 reviewing process would significantly reduce the risk of bias while having minimal impact on approved application quality.

7. Conclusion

We combine state-of-the-art language models and human perception tests to identify plausibly causal USPTO discrimination against under-represented groups and individual inventors in violation of U.S. Code. We are the first to systematically investigate patenting disparities across various dimensions and their interactions/intersectionality. In doing so, we document persistent discrimination towards female and black inventors in patenting approvals. We also identify a novel "affiliation bias" in innovation that favors, in ways unrelated to invention quality, applicants from large and private firm to an extent more significant than demographic discrimination. Bias against individuals based on such affiliation is illegal in the United States, and is unlikely to be unique to patenting. Though such effects are well known in academia, their presence in economic contexts necessitates further study in a wide range of fields. We explain all of our results through a parsimonious theoretical framework based on correlation neglect, which implies discrimination by rational regulators even in the absence of ill intent.

To circumvent the roadblock to socioeconomic mobility and equality, improving the education and living environments for better development of skills in under-represented innovators and entrepreneurs is not sufficient. Government agencies must simultaneously address and rectify unfair discrimination on their part. To that end, we quantify these two forms of important disparities and shed insight on the underlying economic channels. In practice, double-blind patent applications present a simple initial solution to discrimination.

Broadly, our C-HML methodology opens opportunities for a comprehensive study of discrimination in many areas involving decision-makers relying on human perceptions, applications/requests by individuals involving texts, and allocation of scarce resources. In many cases where field experiments are infeasible, a human+machine approach may rigorously and causally distinguish correlations

between demographics and outcomes from human bias, which is critical to identifying solutions and resolving such hitherto irresolvable 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.

Bernstein, S., A. Korteweg, and K. Laws. 2017. Attracting early-stage investors: Evidence from a randomized field experiment. *Journal of Finance* 72:509–38.

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.

Chernenko, S., and D. S. Scharfstein. 2022. Racial disparities in the paycheck protection program. 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.

Chien, C. V., J. Grennan, and J. Sandvik. 2025. Small-scale mentoring, large-scale innovation: Evidence from a superstar firm. Working Paper, Berkeley Law, Emory, UArizona.

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.

Gompers, P., J. Lerner, and D. Scharfstein. 2005. Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999. *Journal of Finance* 60:577–614.

Gompers, P. A., W. Gornall, S. N. Kaplan, and I. A. Strebulaev. 2020. How do venture capitalists make decisions? *Journal of Financial Economics* 135:169–90.

Hacamo, I., and K. Kleiner. 2022. Forced entrepreneurs. *Journal of Finance* 77:49–83.

Hansen, A. L., and S. Kazinnik. 2023. Can ChatGPT decipher fedspeak? Working Paper, Federal Reserve Board.

He, S., L. Lv, A. Manela, and J. Wu. 2025. Chronologically consistent large language models. Working Paper, WashU Olin.

Hoberg, G., and A. Manela. 2025. The natural language of finance. *Available at SSRN 5119322* .

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.

Jha, M., J. Qian, M. Weber, and B. Yang. 2023. ChatGPT and corporate policies. Working Paper, GSU Robinson, Chicago Booth.

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.

Kim, A. G., M. Muhn, and V. V. Nikolaev. 2023. Bloated disclosures: Can ChatGPT help investors process information? Working Paper, Chicago Booth.

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. 2025. Understanding innovation quality and success with large-language models. Working Paper.

Appendix A: Proofs

Proof of Proposition 1.

Let $\beta_2 = 0$. Then $y = \beta_1(q + \epsilon_1) = \beta_1 q_i + \beta_1 q_c + \beta_1 \epsilon_1$. Then

$$corr(y, q_c) = \frac{\text{Cov}(y, q_c)}{\sqrt{\text{Var}(y)\,\text{Var}(q_c)}}. \tag{15}$$

By the independence of q_c from q_i and ϵ_1, we have

$$corr(y, q_c) = \frac{\beta_1\,\text{Cov}(q_c, q_c)}{\sqrt{\text{Var}(y)\,\text{Var}(q_c)}} = \beta_1\sqrt{\frac{\text{Var}(q_c)}{\text{Var}(y)}} = \beta_1\frac{\sigma_{q_c}}{\sigma_y} \tag{16}$$

Clearly, this takes non-zero values if $\beta_1 \neq 0$. Since $\beta_2 = 0$, β_1 must be non-zero as otherwise all applications are identically of 0 quality. Since $corr(y, q_c) \neq 0$, we have $\hat{\beta}_2 \neq 0$ and $\hat{\beta}_2 \neq \beta_2$. □

Proof of Proposition 2.

We first rewrite y in the form

$$y = \beta_1 q + \beta_2 q_c + (\beta_1\epsilon_1 + \beta_2\epsilon_2) = \beta_1(q_{ML}) + \beta_2 q_c + u, \tag{17}$$

$$u = \beta_1\epsilon_1 + \beta_2\epsilon_2 - \beta_1\epsilon_{ML}. \tag{18}$$

Now, let r_{q_c} be the residual obtained from regressing q_c on q_{ML}, i.e.,

$$r_{q_c} = q_c - \gamma q_{ML}, \gamma = \frac{\text{Cov}(q_{ML}, q_c)}{\text{Var}(q_{ML})} \tag{19}$$

Then, by the Frisch-Waugh-Lovell theorem, the expression for $\hat{\beta}_2$ in the estimated regression is given by

$$\hat{\beta}_2 = \frac{\text{Cov}(r_{q_c}, y)}{\text{Var}(r_{q_c})}. \tag{20}$$

Recall that $y = \beta_1 q_{ML} + \beta_2 q_c + u$ and that $\text{Cov}(r_{q_c}, q_{ML}) = 0$ by definition. Then, we can rewrite

$$\hat{\beta}_2 = \frac{\beta_2\,\text{Var}(r_{q_c}) + \text{Cov}(r_{q_c}, u)}{\text{Var}(r_{q_c})} = \beta_2 + \frac{\text{Cov}(r_{q_c}, u)}{\text{Var}(r_{q_c})}. \tag{21}$$

Using the fact that $r_{q_c} = q_c - \gamma q_{ML}$, we have $\text{Cov}(r_{q_c}, u) = \text{Cov}(q_c, u) - \gamma\,\text{Cov}(q_{ML}, u)$. u is a sum

of errors independent from q_c, so

$$\text{Cov}(r_{q_c}, u) = -\gamma \,\text{Cov}\left(q + \epsilon_{ML}, \beta_1(\epsilon_1 - \epsilon_{ML}) + \beta_2\epsilon_2\right) = \gamma\beta_1 \,\text{Var}(\epsilon_{ML}). \tag{22}$$

Additionally, note that

$$\text{Var}(r_{q_c}) = \text{Var}(q_c) - \frac{\text{Cov}(q_{ML}, q_c)^2}{\text{Var}(q_{ML})}. \tag{23}$$

Thus, with some simplification, we have that

$$\hat{\beta}_2 - \beta_2 = \frac{\beta_1 \,\text{Cov}(q_{ML}, q_c)\,\text{Var}(\epsilon_{ML})}{\text{Var}(q_{ML})\,\text{Var}(q_c) - \text{Cov}(q_{ML}, q_c)^2}. \tag{24}$$

Note that $\text{Cov}(q_{ML}, q_c) = \text{Var}(q_c) = \sigma^2_{q_c}$, so

$$\hat{\beta}_2 - \beta_2 = \frac{\beta_1\sigma^2_{q_c}\sigma^2_{ML}}{\sigma^2_{q_{ML}}\sigma^2_{q_c} - \sigma^4_{q_c}} = \beta_1\frac{\sigma^2_{ML}}{\sigma^2_{q_{ML}} - \sigma^2_{q_c}} = \beta_1\frac{\sigma^2_{ML}}{\sigma^2_{q_i} + \sigma^2_{ML}}, \tag{25}$$

the desired expression. □

Proof of Proposition 3.

We will apply the Frisch-Waugh-Lovell theorem to find $\hat{\beta}_2$. Again, let

$$r_{s_N} = s_N - \gamma q_{ML}, \gamma = \frac{\text{Cov}(q_{ML}, s_N)}{\text{Var}(q_{ML})}. \tag{26}$$

Then,

$$\hat{\beta}_2 = \frac{\text{Cov}(r_{s_N}, y)}{\text{Var}(r_{s_N})}. \tag{27}$$

Before proceeding, we recall that $s_N = q_c + \epsilon_N$, so with q_c held constant, ϵ_N is the source of all variation in s_N. Consequently,

$$\text{Cov}(s_N, q_{ML}) = \text{Cov}(q_c + \epsilon_N, q_{ML}) = \text{Cov}(\epsilon_N, q_{ML}) = 0. \tag{28}$$

This is the numerator of the regression coefficient of s_N on q_{ML}, so $\gamma = 0$ and r_{s_N} simply equals s_N.

Returning to the expression for $\hat{\beta}_2$, we have

$$\hat{\beta}_2 = \frac{\text{Cov}(s_N, y)}{\text{Var}(s_N)} = \frac{\beta_1 \,\text{Cov}(s_N, s_1) + \beta_2 \,\text{Var}(s_N)}{\text{Var}(s_N)}. \tag{29}$$

Because q_c is constant, $\beta_1 \,\text{Cov}(s_N, s_1)$ in the numerator is equal to $\beta_1 \,\text{Cov}(\epsilon_N, q_i + \epsilon_1) = 0$. Thus,

$$\hat{\beta}_2 = \beta_2 \frac{\text{Var}(s_N)}{\text{Var}(s_N)} = \beta_2. \tag{30}$$

□

Proof of Proposition 4. Substituting the true model,

$$y = \beta_1 q + \gamma_G \, G + \gamma_R R + \gamma_{GR} \, (G \times R) + \beta_1 \epsilon_1 + \beta_2 \epsilon_N. \tag{31}$$

The term $\beta_1 q$ is completely taken out by q_{ML}, since we assume $q_{ML} = q$. Thus, the remaining terms are given by

$$\eta = \beta_2 \big(\beta_R \, R + \beta_{GR} \, (G \times R) \big) + \beta_1 \epsilon_1 + \beta_2 \epsilon_N. \tag{32}$$

By the omitted variable bias formula, the estimated coefficient $\hat{\beta_G}$ on G is given by:

$$\hat{\gamma}_G = \gamma_G + \frac{\text{Cov} \left(G, \; \gamma_R \, R + \gamma_{GR} \, (G \times R) \right)}{\text{Var}(G)}. \tag{33}$$

Since G is binary (so that $G^2 = G$), we have:

$$\text{Cov}(G, \, G \times R) = \text{Cov}(G, R). \tag{34}$$

Thus, the estimated coefficient becomes:

$$\hat{\gamma}_G = \gamma_G + (\gamma_R + \gamma_{GR}) \, \frac{\text{Cov}(G, R)}{\text{Var}(G)}. \tag{35}$$

The "true" coefficient is γ_G, so the bias is given by

$$\hat{\gamma}_G - \gamma_G = \gamma_R + \gamma_{GR} \frac{\text{Cov}(G, R)}{\text{Var}(G)}. \tag{36}$$

There is no restriction on the potential magnitude of $\gamma_R + \gamma_{GR}$, so $\hat{\gamma}_G$ could be over-estimated, under-estimated, or of a different sign than the desired γ_G. □

Proof of Proposition 5.

Because all four normal variables $(q_i, q_c, \epsilon_1, \epsilon_N)$ are independent and of zero mean, the system $(q, s_1, s_N)^T$ follows a multivariate normal distribution. Therefore, the conditional expectation of q on $(s_1, s_N)^T$ is given by

$$E[q|s_1, s_N] = \text{Cov}[q, (s_1, s_N)]\text{Var}[(s_1, s_N)^T]^{-1} \cdot (s_1, s_N)^T. \tag{37}$$

Again using the independence of $(q_i, q_c, \epsilon_1, \epsilon_N)$, we solve

$$\text{Cov}[q, (s_1, s_N)] = (\sigma_i^2 + \sigma_c^2, \sigma_c^2). \tag{38}$$

Similarly,

$$\text{Var}[(s_1, s_N)^T] = \begin{bmatrix} \sigma_i^2 + \sigma_c^2 + \sigma_{\epsilon_1}^2 & \sigma_c^2 \\ \sigma_c^2 & \sigma_c^2 + \sigma_{\epsilon_N}^2 \end{bmatrix}. \tag{39}$$

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

$$\text{Cov}[q, (s_1, s_N)] \cdot \text{Var}[(s_1, s_N)^T]^{-1} = \left(\frac{\sigma_i^2 + \frac{\sigma_c^2 \sigma_{\epsilon_N}^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2}}{\sigma_i^2 + \frac{\sigma_c^2 \sigma_{\epsilon_N}^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2} + \sigma_{\epsilon_1}^2}, \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2} \cdot \frac{\sigma_{\epsilon_q}^2}{\sigma_i^2 + \frac{\sigma_c^2 \sigma_{\epsilon_N}^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2} + \sigma_{\epsilon_1}^2} \right).$$

For ease of notation, we write

$$\beta_1 = \frac{\sigma^2}{\sigma^2 + \sigma_{\epsilon_q}^2}, \quad \sigma^2 = \sigma_i^2 + \frac{\sigma_c^2 \sigma_{\epsilon_N}^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2}$$

such that

$$E[q|s_1, s_N] = \left(\beta_1, \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2} \cdot (1 - \beta_1) \right) \cdot (s_1, s_N)^T. \tag{40}$$

□

Proof of Proposition 6.

Similarly to the proof of Proposition 5, the conditional expectation of q on $(s_1', s_N)^T$ is given by

$$E[q|s_1', s_N] = \text{Cov}[q, (s_1', s_N)]\text{Var}[(s_1', s_N)^T]^{-1} \cdot (s_1', s_N)^T. \tag{41}$$

The expressions for both matrices are simpler in this case:

$$\text{Cov}[q, (s_1', s_N)] = (\sigma_i^2, \sigma_c^2) \tag{42}$$

$$\text{Var}[(s_1', s_N)^T] = \begin{bmatrix} \sigma_i^2 + \sigma_{\epsilon_q}^2 & 0 \\ 0 & \sigma_c^2 + \sigma_{\epsilon_N}^2 \end{bmatrix}. \tag{43}$$

In the same manner as before, we solve

$$\text{Cov}[q, (s_1', s_N)] \cdot \text{Var}[(s_1', s_N)^T]^{-1} = \left(\frac{\sigma_i^2}{\sigma_i^2 + \sigma_{\epsilon_1}^2}, \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2} \right), \tag{44}$$

such that

$$E[q|s_1', s_N] = \left(\frac{\sigma_i^2}{\sigma_i^2 + \sigma_{\epsilon_1}^2}, \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2} \right) \cdot (s_1', s_N)^T. \tag{45}$$

Note that $\beta_2' = \beta_2 \cdot \dfrac{1}{1-\beta_1}$. Since $\beta_1 \in (0,1)$, $\beta_2' > \beta_2$. Meanwhile, $\beta_1 = \dfrac{\sigma^2}{\sigma^2 + \sigma_{\epsilon_1}^2}$ and $\beta_1' = \dfrac{\sigma^2 - x^2}{\sigma^2 + \sigma_{\epsilon_1}^2 - x^2}$, where $x^2 = \dfrac{\sigma_c^2 \sigma_{\epsilon_N}^2}{\sigma_c^2 + \sigma_{\epsilon_N}^2} < \sigma^2$, and thus $\beta_1' < \beta_1$. □

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. 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 4. 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 5. Discrimination and Gendered Names

This table reports the results of the regressions of application acceptance on perceived inventor gender, sorted into two groups by actual 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)
Perceived Male-ness	0.076***	0.076***	0.077***	0.017***	0.017***	0.018***
	(26.02)	(18.83)	(19.16)	(6.27)	(6.30)	(6.53)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	129,978	129,978	129,978	129,978	129,978	129,978
Adjusted R^2	0.005	0.005	0.018	0.123	0.123	0.127

Panel B: Male Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Perceived Male-ness	0.053***	0.053***	0.059***	0.008***	0.008***	0.011***
	(26.78)	(12.86)	(15.81)	(4.23)	(3.62)	(5.39)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	1,405,341	1,405,341	1,405,341	1,405,341	1,405,341	1,405,341
Adjusted R^2	0.001	0.001	0.013	0.108	0.108	0.111

Table 6. Discrimination and Gender-Ambiguous 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 7. Preference for Whites and Name Perception

This table reports the results of the regressions of application acceptance on perceived name white-ness, sorted into two groups by actual race (white vs. non-white). 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: True White Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Perceived White-ness	0.009***	0.009***	0.009***	0.005***	0.005**	0.005***
	(7.39)	(4.27)	(4.74)	(4.70)	(2.85)	(3.05)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	669,756	669,756	669,756	669,756	669,756	669,756
Adjusted R^2	0.000	0.000	0.012	0.107	0.107	0.110

Panel B: Non-White Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Perceived White-ness	0.025***	0.025***	0.031***	0.019***	0.019***	0.022***
	(9.08)	(3.91)	(5.48)	(7.56)	(4.36)	(5.18)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	530,584	530,584	530,584	530,584	530,584	530,584
Adjusted R-squared	0.000	0.000	0.014	0.105	0.105	0.108

Table 8. Discrimination against Black Inventors and Name Perception

This table reports the results of the regressions of application acceptance on perceived name black-ness, sorted into two groups by actual race (black vs. non-black). 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: True Black Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Perceived Black-ness	-0.043***	-0.043***	-0.046***	-0.031***	-0.031***	-0.033***
	(-5.50)	(-5.05)	(-5.40)	(-4.19)	(-4.48)	(-4.79)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	17,633	17,633	17,633	17,633	17,633	17,633
Adjusted R^2	0.002	0.002	0.015	0.130	0.130	0.134

Panel B: Non-Black Inventors.

	Acceptances					
	(1)	(2)	(3)	(4)	(5)	(6)
Perceived Black-ness	-0.010***	-0.010**	-0.003	-0.005***	-0.005*	-0.002
	(-5.40)	(-2.48)	(-1.11)	(-2.81)	(-1.91)	(-1.04)
Year-Fixed Effects	-	-	Y	-	-	Y
Clustered by Year	-	Y	Y	-	Y	Y
Application Quality Control	-	-	-	Y	Y	Y
N	1,182,707	1,182,707	1,182,707	1,182,707	1,182,707	1,182,707
Adjusted R^2	0.000	0.000	0.012	0.105	0.105	0.108

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. 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 12. 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

Table 13. 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 14. 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 15. 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 16. 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 17. 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 18. 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