Implicit Quotas

Roland G. Fryer, Jr.*

Harvard University and NBER

Abstract

Employment or admission "goals" are often preferred to affirmative action as a way of obtaining diversity. By constructing a simple model of employer-auditor interaction, it is shown that when an auditor has imperfect information regarding employers' proclivities to discriminate and the fraction of qualified minorities in each employers applicant pool, goals are synonymous with quotas. Technically speaking, any equilibrium of the auditing game involves a non-empty set of employers that hire so that they do not trigger an audit by rejecting qualified non-minorities, hiring unqualified minorities, or both. Further, under some assumptions, explicit quotas (those mandated by an auditor) are more efficient than implicit quotas (goals settled upon in equilibrium by employers wishing to avoid an audit).

*Fryer is a Professor of Economics at Harvard Univeristy and Faculty Research Fellow at the National Bureau of Economic Research (NBER). Correspondence can be addressed to: Department of Economics, Harvard University; 1805 Cambridge Street; Cambridge, MA 02138. E-mail: rfryer@fas.harvard.edu.

I would like to thank Lavone Norwood, Corporate Executive Director of Employee Relations, for extensive discussions on auditing processes, Mary Leone, Director of Equal Employment Opportunity/Affirmative Action compliance, and several anonymous employers for helpful suggestions. I would also like to thank Gary Becker, Kalyan Chatterjee, Tom Gresik, Nezih Guner, James Heckman, Tony Kwasnica, Steve Levitt, Glenn Loury, Phil Reny, David Shapiro, Tomas Sjöström, Lars Stole, and seminar participants too numerous to mention. Jörg Spenkuch provided exceptional research assistance. A portion of this research was conducted while Fryer was a research assistant in the Equal Employment Opportunities Department of an anonymous corporation. The usual disclaimer applies.

"Since President Nixon was here in my job, America has used goals and timetables to preserve opportunity and to prevent discrimination, to urge businesses to set higher expectations for themselves and to realize those expectations. But we did not and we will not use rigid quotas to mandate outcomes."

— *President William J. Clinton, July 19, 1995.*

"I am for Affirmative Action, as I describe it, but not for quotas or preferences."

— *President George W. Bush, April 2, 2000.*

"We do not think it matters whether a government hiring program imposes hard quotas, soft quotas, or goals. Any of these techniques induces an employer to hire with an eye toward meeting the numerical target.""

— *Judge Laurence Silberman, Lutheran Church–Missouri Synod v. FCC*

1. INTRODUCTION

In its 40-year history, there have been many popular misconceptions about affirmative action, including: (1) the only way to create a color-blind society is to adopt color-blind means; (2) Affirmative Action may have been necessary 30 years ago, but the "playing field" is fairly level now; (3) The public doesn't support affirmative action anymore; (4) a large percentage of white workers will lose out if affirmative action is continued; and (5) goals and timelines are better than rigid affirmative action quotas. The first four misconceptions have been shown to be more myth than fact (Fryer and Loury 2005; Fryer, Loury, and Yuret forthcoming; Loury 1979). Yet, goals are thought to be good faith efforts on the part of noble employers, whereas quotas are envisioned as rigid racial diversity requirements that often result in the hiring of incompetent minorities. In its landmark decision *Regents of the University of California v. Bakke* (438 U.S. 265) the Supreme Court ruled such inflexible quotas to be unconstitutional; while it upheld the use of soft quotas or goals in *Fullilove v. Klutznick* (448 U.S. 448). Understanding the relationship between these amorphous terms is the subject of this paper.

To get beneath the terminology, I develop a model of employer-auditor interaction that involves imperfect auditing of an employer's hiring practices. Employers differ in their proclivities to discriminate and in the fraction of qualified minorities that apply for positions in their firm. After observing their type (desire to discriminate and applicant pool), each employer hires a ratio of minorities to non-minorities. Thus, if an employer hires a small share of minorities it implies one of two things: either the employer is a discriminator who rejected some qualified minority candidates or a non-discriminator who had a small fraction of qualified minorities apply. I assume that an outside auditor cannot distinguish perfectly between these states – even after an audit. The auditor observes each employers workforce and decides whether or not to conduct an audit. So, in an effort to eliminate discrimination, the auditor will mistakenly punish employers who did not discriminate while others (who actually did discriminate) go undetected. It is this informational asymmetry that gives employers incentives to alter their hiring ratio so as not to induce an audit.

The results of the simple auditing model are illuminating: all equilibria exhibit an implicit quota property. That is, a non-empty set of employers (both those who are inclined to discriminate, and those who are not) are willing to alter their behavior to avoid an audit, since there is a positive probability that the auditor makes a mistake and the penalty is strictly positive. I use the modifier "implicit" for a particular reason. If we were deriving *explicit* quotas, this would be represented by the government announcing a desired ratio of minorities to non-minorities, and (assuming the penalty for deviating is sufficiently large) employers strictly adhering to this ratio, which has been ruled unconstitutional (*Regents of the University of California v. Bakke* (438 U.S. 265)). In contrast, implicit quotas are those that the employers themselves set, in equilibrium, as an optimal response to imperfect auditing. Thus, the quotes that we began with and the rhetoric from both political parties that supports goals but not quotas has no content.

The lesson is straightforward. If a regulator is interested in enforcing anti-discrimination laws, then *goals are quotas when an auditor has imperfect information regarding employers' desires to discriminate and the fraction of qualified minorities who applied to each firm.* Under some assumptions, explicit quotas are more efficient than implicit quotas.[1]

This model, although applied here to auditing in the labor market (where quotas are most controversial), can naturally be applied to auditing environments, discusses the equilibrium of the model, involving tax evasion, teacher accountability, and anti-trust enforcement.[2]

[1] Imperfect information is the crucial assumption. Without it, goals and quotas can be quite different objects. But, in practice, the information auditors have on an employers discriminatory intentions or the quality of their applicant pools is far from perfect.

[2] This paper is related to the well developed literatures on employment discrimination and tax compliance. There

The paper proceeds as follows: Section 2 provides a concise, but relatively informal, verbal description of an auditing model with imperfect information, and constructs a numerical example which illustrates the main results. Section 3 concludes. Appendix A contains the formal model along with technical proofs of all the results discussed in Section 2. Appendix B provides additional results from the model.

2. A MODEL OF IMPERFECT AUDITING

Let there be a continuum of workers and a continuum of employers. Workers belong to one of two groups: "minorities" or "non-minorities." There are also two types of employers: some are "biased" against minorities while the others are "unbiased." There is a set of auditors. Before the start of the game, the Government chooses a penalty to be enforced on employers who discriminate against minorities in their hiring practices.

Nature moves first and assigns a two-dimensional type to each employer: whether or not they are biased and the quality of their applicant pool. The latter is a number on the positive real line, distributed according to a smooth and continuous cumulative distribution function. One can think of this number as the profit maximizing (absent discriminatory taste) ratio of minorities to non-minorities in an employer's applicant pool.[3] This formulation is flexible enough to allow for

is a relatively large literature on employment discrimination. The two main theories are Becker (1957) and Arrow (1973). Becker (1957) provides a taste-based theory of discrimination. In this theory, agents discriminate because there exists non-pecuniary psychic costs to interaction with minorities. Thus, in this model, agents are willing to forgo profits or earn lower wages to ensure segregation. Arrow (1973) discusses a model of statistical discrimination. This model shows that employers can (rationally) discriminate against a group even when they are ex ante identical. Independent of the underlying theory of discrimination, it manifests itself in our model by an employer hiring fewer minorities than they otherwise would. A regulator, then, may want to break such equilibria. This auditing problem is similar to the extensive literature on tax compliance, for which there is an impressive literature (see Andreoni, Erard, and Feinstein [1998], for an overview). Surprisingly, however, the literature on imperfect auditing, is small (Lawarree and Van Audenrode 1992; Kofman and Lawarree 1993; Bardsley 1996; and implicitly, Erard and Feinstein 1994). The closest to my approach is Bardsley (1996) who has a similar model of imperfect auditing. The key difference involves the auditors payoffs.

[3]Ideally, one would want to endogenize the employer's state and allow employers to make investments to increase their likelihood of being in a "good" state, using the monotone likelihood ratio property. I do not model these initial investments by the employers since they are not observable by the auditor. If it helps to fix ideas, one can assume that the "lottery" of states is determined by investment (i.e. recruitment) activities of firms outside of our model. However, if the function that maps recruiting initiatives into applicant pools is not deterministic (i.e. intense minority

different distributions of effort, investment, talent, geography, or other factor that might change the profit maximizing ratio of minorities to non-minorities an employer wants to hire.

Next, employers observe their two dimensional private type and make hiring decisions. The following provides a formal definition of discrimination.

Definition 1 *A firm is said to discriminate if it hires a ratio of minorities to non-minorities that is strictly less than its profit maximizing ratio (not including their possible discriminatory taste).*

By definition, unbiased employers hire the profit-maximizing ratio of minorities to non-minorities, absent regulation. I further assume that biased employers, absent regulation, hire strictly less. In the language of Becker (1957), one can think of this difference as a discrimination coefficient or in statistical or cognitive discrimination models, it may capture the lower share of blacks hired due to negative stereotypes (Arrow 1973) or coarser categories (Fryer and Jackson forthcoming). An auditor, after observing each employer's hiring decision (not their type or applicant pool), makes a dichotomous audit decision: audit or not. If the auditor decides to conduct an audit, she makes a correct assessment of the employer with probability greater than a half and with the complementary probability she makes a mistake. After the audit, the auditor decides whether to issue the fine or not.[4] It is important to emphasize that discrimination here is thought to be one-sided: a regulator is auditing hiring practices to lessen discrimination against minorities which exists absent regulation. In a more elaborate model, one can add penalties for overshooting and discriminating against non-minorities.

Employer's payoffs are represented by a single peaked function which reaches its maximum when they hire their optimal ratio of minorities to non-minorities – taking into account their applicant pool and possibly discriminatory preferences. Auditors receive a positive payoff if they punish a discriminating employer and suffer a cost if they do not fine a discriminating employer or mistakenly fine a non-discriminator.[5] The payoff to an auditor for not punishing a non-discriminator

recruiting need not always result in a minority rich applicant pool), then the exogenous determination of states is without further loss of generality.

[4] Realistically, the punisihment should be proportional to the level of discrimination. Adding more elaborate penalty functions will make interesting changes in the qualitative properties of the equilibria. But, if the penalty is strictly positive for all discriminatory acts – our main result holds.

[5] Assuming that the auditor receives negative utility from punishing a non-discriminator is equivalent to the limited liability assumption in the tax literature.

is normalized to zero.

It is assumed that auditors are only interested in finding and punishing employers who discriminate against minorities. She does not care about their (possibly) biased preferences towards minorities, as long as they do not use discriminatory hiring practices. To keep things simple, I assume that all payoff relevant parameters are exogenously given.

2.1. Equilibrium

To solve the model, I focus on pure strategy equilibria in which each agent makes a deterministic choice and all individuals of the same type make the same choice.[6]The solution concept for the auditing model, per usual for signaling games, is perfect Bayesian equilibrium. Intuitively, a perfect Bayesian equilibrium is a set of strategies and beliefs such that, at any stage of the game, strategies are optimal given the beliefs, and beliefs are obtained from equilibrium strategies and observed actions using Bayes's rule. In what follows, I describe existential results for two possible sets of equilibria: separating and semi-separating. All formal statements of the propositions, along with their proofs, can be found in Appendix A. Appendix B treats the possibility of pooling, non-monotonic, and mixed strategy equilibria.

2.1.1. Separating Equilibrium

For (standard) monotonic signaling models, in a fully separating equilibrium, an agent of each type chooses a unique action, and each type is correctly identified in equilibrium. In this (slightly non-standard) non-monotonic signaling game, full separation is ruled out *a priori*, due to continuous types and strategy spaces. In particular, for every hiring ratio there exists two applicant pools – one richer in qualified minorities than the other – such that a biased employer with the better pool will hire the same number of minorities as an unbiased employer with less qualified minorities among her applicants. Consider the following (slightly perturbed) definition of a separating equilibrium.

Definition 2 *In any pseudo-separating equilibrium, each employer hires their profit maximizing*

[6]A strategy for an employer is an assignment function that maps their private type into a ratio of minorities to non-minorities hired. A strategy for an auditor is a function that maps an employers observed ratio of minorities to non-minorities hires into an audit decision.

To begin, I restrict attention to monotonic strategies for the auditor (i.e. "cutoffs"), in which she audits any employer with observed minority:non-minority ratio less than the cutoff and does not audit any employer with an observed ratio above it. This is without loss of generality when the auditor uses pure strategies (see Appendix B).

ratio of minorities to non-minorities

Our first result (Proposition 1, Appendix A) highlights the fact that no pseudo-separating equilibrium exists. Any effort on the auditor's part to find and punish discriminators will necessarily yield an implicit quota (hire more minorities than one would in their profit maximizing workforce). This is the main theme of the paper. The surprising part of this result is that biased and unbiased employers alike may adhere to the implicit quota.

Some believe that quota-like hiring from employers would be an easier pill to swallow if the market was not accounting for tastes, and the threat of an audit simply forced biased employers to hire workforce ratios that were equivalent to what unbiased employers would optimally hire, conditional on the same applicant pool.[7] This is possible if the auditor has perfect information. However, some may find it disturbing that, given the auditors' lack of information, even unbiased employers are willing to alter their hiring ratios as not to induce an audit. Especially when (as Proposition 1 proves) this behavior is inevitable, for at least some employers.

2.1.1. Semi-Separating Equilibria

There are two types of semi-separating equilibria, which I label "marginal" and "inclusive." The distinction between them turns on what types of employers choose to pool on the implicit quota. Marginal equilibria require that only marginal employers (employers whose profit maximizing hiring ratios are relatively close to the implicit quota) adhere to the implicit quota. In this type of equilibrium, employers with applicant pools that have very few qualified minorities refuse to alter their hiring ratios enough to avoid an audit, because the profit loss in doing so is large relative to the expected cost of being audited. They simply incur the expected cost. As the fine for being deemed a discriminator gets large, fewer employers will risk the penalty and inclusive equilbria obtain. In an inclusive semi-separating equilibrium all employers (whether or not they are biased), who face a profit maximizing hiring ratio below the implicit quota, will alter their behavior and hire right up to the implicit quota. Notice: marginal equilibria are the only equilibria for which audits occur in equilibrium. In this sense, one may find them more appealing and empirically relevant.

Whether marginal or inclusive equilibria obtain depends solely on the magnitude of the penalty

[7] However, others believe that not accounting for market tastes is a mistake, even if if means that some groups endure discriminatory treatment (see Epstein [1994]).

for discriminating. If the expected penalties are relatively small marginal equilibria exist (Proposition 2, Appendix A); if the expected penalties are large inclusive equilibria obtain (Proposition 3, Appendix A).

The technical conditions to ensure a marginal equilibrium require that all employers with relatively small fractions of qualified minorities in their applicant pools hire their profit maximizing workforce. This puts an upper bound on penlalty that can be imposed in equilibrium. The conditions also ensure that there is always a set of positive measure of employers who alter their profit maximizing workforce by "jumping up" to the implicit quota And, given this behavior from employers, the auditor does not find it worthwhile to audit at that quota. The conditions for inclusive equilibrium ensure that the expected costs of being audited are sufficiently high to dissuade potential deviators, and the implicit quota is high enough to minimize the amount of discrimination in equilibrium, so that the auditor does not find it optimal to audit the employers that pile up at the quota even though she knows that some of them are discriminating.

A simple numerical example illustrates many of the points stressed thus far. Assume the auditor believes an employer who hires a certain ratio of non-minorities to minorities to be a discriminator with probability of one minus the hiring ratio. Further, assume the auditor receives a payoff of 1 if she correctly punishes a discriminator, and incurs a penalty of 1 if she mistakenly fines a non-discriminating employer. Finally, assume the probabilty of making a correct assessment is 80%, and the costs of conducting an audit are fixed at $\frac{7}{10}$. Under these circumstances the auditor will find it optimal to audit all employers who hires less than 50% minorities (see Appendix A for derivation). The employer utility function obtains a maximum of 0 when they hire their profit maximizing ratio. Whenever an employer deviates from this ratio he receives a penalty equal to the squared deviaton. Also, suppose discriminating employers will hire only half the number of minorities an otherwise equal unbiased employer would hire. Under these assumptions a marginal semi-seperating equilibrium will exist if the fine is set to $\frac{5}{64}$. Then, all employers, whether they are discriminators or not, with optimal hiring ratios between $\frac{3}{8}$ and $\frac{1}{2}$ will alter their behavior and hire 50% minorities. Those employers who would hire less than $\frac{3}{8}$ will not change their hiring decisions, but rather incur the expected costs of being audited. Employers who would hire more than 50% anyway will also not deviate, as they will not be audited in equilibrium. If, however, the fine is set higher than $\frac{5}{4}$, then an inclusive semi-separating equilibrium exists. That is all employers who would otherwise hire a fraction of minorites smaller than $\frac{1}{2}$ will now hire exactly right up to this line to avoid the possibility of being fined by the auditor. All other employers, will again, not deviate from their optimal hiring ratio.

2.2. The Multiplicity Problem

In typical signaling models, one is plagued with the multiplicity of equilibria due to the freedom associated with out of equilibrium beliefs in standard solution concepts. For example, suppose we have an inclusive equilibrium with a lower bound hiring ratio, and the auditor happens to observe an employer who hires below that ratio. In this case, perfect Bayesian does not specify the auditors inferences, thus, it is theoretically plausible that an auditor will believe that any deviations below are certainly discriminators. He could just as easily believe that they are non-discriminators. He is free to choose. As a result of the lacksadaisical requirements on out of equilibrium beliefs imposed by Bayesian perfection, we have a continuum of potential equilibrium (Proposition 4 and Proposition 5, Appendix A). For instance, we know that in any equilibrium the auditor does not want to audit employers who pile up on the implicit quota. Well, there are a continuum of possible implicit quotas above that which makes an auditor indifferent. With out of equilibrium beliefs that anyone who hires beneath the implicit quota is a discriminator – all of these possibilities are equilibria.

This type of multiplicity problem is an unfortunate result that stifles the predictive power of most signaling models. However, it can be argued that the out of equilibrium beliefs needed to construct the equilibria above are not empirically relevant. In particular, it may be unreasonable to assume that every deviation from a candidate equilibrium is a discriminator. Cho and Kreps (1986) posed an equilibrium refinement known as the intuitive criterion. This criterion was constructed to aid in choosing between the multiplicity of possible equilibria found in most signaling games. The criterion is applied in our model in a series of steps.

1. For any deviation out of a candidate equilibrium, define a set of types that would get less than their equilibrium payoff by making the deviation, provided the auditor played an undominated strategy.

2. Define a set of types that would necessarily be better off by employing the deviating ratio, given the auditor knows that the deviation could not have come from any employer in the set

defined in step 1.

3. If the set defined in step 2 is non empty, the equilibrium fails the intuitive criterion.

In many (standard) signaling models, this refinement has eliminated the multiplicity problem. Cho and Kreps (1986) show, in the Spence model of job market signaling (Spence 1974) that the intuitive criterion selects the separating equilibrium with the least amount of inefficient signaling. Unfortunately, it has absolutely no bite in the current (non standard) model.[8] Proposition 6 shows that all semi-separating equilibria (marginal and inclusive) survive after applying the intuitive criterion.

To see this, consider the three step verification process outlined above. The proposition shows that the set of employers who strictly prefer their equilibrium payoff to any deviation is precisely the set of employers that are hiring their first best above the implicit quota. We know that an equilibrium fails the intuitive criterion if there exists a type who would be necessarily better off by deviating given the auditor will know that the deviation did not come from any employer hiring their first best. However, there does not exist such a type, because there is still a positive probability being penalized even when you are not discriminating. In other words, imperfect information after the audit undermines the Intuitive Criterion.

3. CONCLUDING REMARKS

Many individuals have an allergic reaction to the use of quotas, but seemingly want to eliminate discrimination by enforcing antidiscrimination laws. *The main result in this paper shows that enforcing anti-discrimination policy has the unintended effect of causing all equilibria to involve a set of employers who alter their hiring ratios to avoid being audited, due to the auditor's lack of information.* In essence, goals are quotas whenever auditing technology is not perfect. And, under some assumptions, goals and targets can lead to more extreme quota-like hiring. Attacking affirmative action as a "quota for minorities" while endorsing "goals" and anti-discrimination enforcement is vacuous.

These results extend in natural and interesting ways to other realms of law and economics. For example, in a tax evasion model, the results indicate that there exists equilibrium in which "honest" tax payers are willing to over report so that they are not fined. Future research in these areas can be

[8] This is also true for the stronger D1 and Equilibrium Dominance Test.

extended along many dimensions. First, it would be interesting to construct a dynamic or repeated model of the auditing process in order to highlight the difficulties an auditor has in identifying discrimination in promotion policies relative to discrimination in initial hiring. From conversations with auditors, they indicate that the former is much more difficult to monitor. A repeated game model would have the advantage that employers could build reputations. Second, as mentioned in the text, one may want to endogenize the applicant pools in two dimensions: (1) allow workers to make human capital investments, and (2) allow the employers to invest in recruiting initiatives in hopes to be given a better pool of potential workers. Another viable extension might be a model in which one allows workers to accuse their employer of discriminating. In fact, the Department of Labor has a discrimination complaint form on their web page. Since most of the money collected by the OFCCP is distributed among those workers who filed the complaints, it may be interesting to examine the strategic relationships at play within this environment.

References

Andreoni, James, Brian Erard, and Jonathan Feinstein. 1998. Tax Compliance. *Journal of Economic Literature* 36:818–60.

Arrow, Kenneth J. 1973. The Theory of Discrimination. Pp. 3–33 in *Discrimination in Labor Markets,* edited by Orley Ashenfelter and Albert Rees. Princeton, NJ: Princeton University Press.

Bardsley, Peter. 1996. Tax Compliance Games With Imperfect Auditing. *Public Finance* 51:473–489.

Becker, Gary S. 1957. *The Economics of Discrimination.* Chicago: University of Chicago Press.

Cho, In-Koo, and David M. Kreps. 1987. Signaling Games and Stable Equilibria. *Quarterly Journal of Economics* 102:179–222.

Epstein, Richard A. 1994. *Forbidden Grounds: The Case Against Antidiscrimination Laws.* Cambridge, MA: Harvard University Press.

Erard, Brian, and Jonathan Feinstein. 1994. Honesty and Evasion in the Tax Compliance Game. *Rand Journal of Economics* 25:1–19.

Fryer, Roland G., and Glenn C. Loury. 2005. Affirmative Action and Its Mythology. *Journal of Economic Perspectives* 19:147–62.

Fryer, Roland G., Glenn C. Loury, and Tolga Yuret. forthcoming. An Economic Analysis of Color-Blind Affirmative Action. *Journal of Law, Economics and Organization.*

Fryer, Roland G., and Matthew O. Jackson. forthcoming. A Categorical Model of Cognition and Biased Decision-Making. *Contributions in Theoretical Economics. Berkeley Electronic Press*

Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory.* Cambridge, MA: MIT Press.

Kofman, Fred, and Jaques Lawarree. 1993. Collusion in Hierarchal Agency.*Econometrica* 61:629–656.

Lawarree, Jaques, and Mark Van Audenrode 1992. Cost Observation, Auditing, and Limited Liability. *Economic Letters* 39:419–423.

Loury, Glenn C. 1995. *One by One from the Inside Out: Essays and Reviews on Race and Responsibility in America.* New York: Free Press.

Phelps, Edmund S. 1972. The Statistical Theory of Racism and Sexism. *American Economic Review* 62:659–61.

Spence, A. Michael. 1974. Job Market Signaling. *Quarterly Journal of Economics* 87:355–374.

APPENDIX A: FORMAL MODEL AND PROOFS OF PROPOSITIONS

A. The Basic Building Blocks

There is a continuum of workers and a continuum of employers, each with unit measure. There are two groups of workers: a measure λ are "minorities" and $1-\lambda$ are "non-minorities." There are also two types of employers: a measure μ are "biased" against minorities and a measure $1-\mu$ are "unbiased." There is also a large set of auditors. The Government chooses a fine P before the start of the game.

Nature moves first and assigns a type (t, a) to each employer, where $t = b$ (resp. $t = u$) if an employer is biased (resp. unbiased), and distributes an applicant pool $a \in [, \overline{a}]$ to each employer according to a smooth and continuous cumulative distribution function $F(a)$ and related density $f(a)$, where a represents the profit maximizing (absent discriminatory taste) ratio of minorities to non-minorities in an employers applicant pool. To avoid trivialities, we assume that every applicant pool has at least one black candidate. Next, employers observe their two dimensional private type, (t, a), and make a workforce decision $r(t, a) \in [0, \infty)$.

By definition, $r(u, a) = a$, that is, unbiased employers hire the profit-maximizing ratio of minorities to non-minorities, absent regulation. When optimizing, biased employers will hire a ratio, $(1-\alpha)a$, $\alpha \in (0, 1)$, absent regulation. An auditor, after observing r (not t or a), makes a dichotomous audit decision: audit or not. If the auditor decides to conduct an audit, she makes a correct assessment of the employer with probability $\phi > \frac{1}{2}$ and with probability $1-\phi$ she makes a mistake. After the audit, the auditor makes a punishment decision; deciding whether to issue the fine P or not.

B. Payoffs

Employer's payoffs are represented by a function $\Gamma(r(u,a) - a)$ for unbiased employers and a function $\Gamma(r(b,a) - (1-\alpha)\,a)$ for biased employers. I make the following assumptions on $\Gamma(\cdot)$.

Assumption 1 *The function $\Gamma(z)$ is twice continuously differentiable, strictly concave, symmetric ($\Gamma(z) = \Gamma(-z)$), and achieves a maximum of zero when $z = 0$.*

Let $\beta > 0$ denote the costs to the auditor of engaging in an audit of an employer. Auditors receive a payoff $\chi > \beta$ if they punish a discriminating employer and suffer a cost $-c < 0$ if they do not fine a discriminating employer or mistakenly fine a non-discriminator. The payoff to an auditor for not punishing a non-discriminator is normalized to zero. I assume that P, χ, c, and β, are exogenously given.

C. Strategies

A strategy for an employer is an assignment function that maps their private type (t, a) into a ratio of minorities to non-minorities hired. A strategy for an auditor is a function that maps an employers observed ratio of minorities to non-minorities hires into an audit decision.

To begin, I restrict attention to monotonic strategies (Q) for the auditor (i.e. "cutoffs"), in which she audits any employer with observed minority:non-minority ratio $r < Q$ and does not audit any employer with an observed ratio $r \geq Q$. This is without loss of generality when the auditor uses pure strategies (see Appendix B).

D. Expected Payoffs

Let $\Psi(r)$ denote the probability that the employer is discriminating, conditional on hiring a workforce r.[9] The auditor's expected payoff of not conducting an audit is $-\Psi(r)c$. When optimizing, with probability $\Psi(r)\phi + (1-\Psi(r))\,[1-\phi]$ the auditor believes that the employer is discriminating, and punishes him with payoff $\Psi(r)\phi\chi - (1-\Psi(r))\,[1-\phi]\,c$. With probability $\Psi(r)\,[1-\phi] + (1-\Psi(r))\,\phi$ she thinks that the employer is not discriminating, and does not punish him (since she receives negative payoff for doing so). The auditor's expected payoff can be written as:

$$\Psi(r)\phi\chi - (1-\phi)\,c - \beta$$

[9] The explicit derivation of $\Psi(r)$ will be equilibrium specific.

The employer's expected payoffs depend on his profit maximizing hiring ratio, the auditor's cut-off strategy, and the expected cost of being audited. An unbiased employer's expected payoff of employing a workforce $r(u,a)$ is $\Gamma(r(u,a)-a)-\phi P\delta_{r<Q}$ if he discriminates, and $\Gamma(r(u,a)-a)-(1-\phi)P\delta_{r<Q}$ if he does not**,** where δ is a standard indicator function**.** Similarly, a biased employer's expected payoff of employing a workforce $r(b,a)$ is $\Gamma(r(b,a)-(1-\alpha)a)-\phi P\delta_{r<Q}$ if he discriminates, and $\Gamma(r(b,s)-(1-\alpha)a)-(1-\phi)P\delta_{r<Q}$ if he does not.

With this notation in hand, we can provide the values needed to recreate the numerical example discussed in the text: $\Psi(r)=1-r;\quad \phi=\frac{4}{5};\quad \chi=c=1;\quad \beta=\frac{7}{10}$. The resulting Q is $\frac{1}{2}$. $\Gamma(r)=-(a-r)^2$ and $\Gamma(r)=-\left((1-\alpha)a-r\right)^2$ for unbiased and biased employers, respectively; $\alpha=\frac{1}{2};\quad P=\frac{5}{64}$ for a marginal equilibrium and $P>\frac{5}{4}$ for an inclusive equilibrium; $\hat{r}=\frac{3}{8}$

Proposition 1 *No pseudo-separating equilibrium exists.*

Proof. To see that no pseudo-separating equilibrium exists, it is sufficient to show that at least one employer will have an incentive to deviate from their first best whenever $\phi<1$. Consider the employer who has a profit maximizing workforce slightly less than the auditor's threshold (Q). (We know this employer exists because of the continuity assumptions.) For an unbiased employer whose first best is slightly less than the expected auditing threshold, the following equation must hold:

$$\Gamma(0)-(1-\phi)P>\Gamma(\varepsilon)$$

$\Gamma(0)=0$ by assumption, so

$$-(1-\phi)P>\Gamma(\varepsilon).$$

However, for any fixed $\phi P>0$, there exists an ε small enough such that this inequality does not hold. Note: $Q=0$ is ruled out by the definition of r^*. *Q.E.D.* ■

Next, I provide 2 definitions.

Definition 3 *In a marginal semi-separating equilibrium,* $r(u,a)=Q$ *for all* $a\in[\widehat{a},Q]$ *and* $r(u,a)=a$ *for all* $a\in[0,\widehat{a}]\cup[Q,\overline{a}]$*; and* $r\left(b,\frac{a}{1-\alpha}\right)=Q$ *for all* $a\in\left[\frac{\widehat{a}}{1-\alpha},\frac{Q}{1-\alpha}\right]$ *and* $r\left(b,\frac{a}{1-\alpha}\right)=a$ *for all* $a\in[0,\widehat{a}]\cup[Q,\overline{a}]$.

Definition 4 *In an inclusive semi-separating equilibrium,* $r(u,a)=Q$ *for all* $a\in[0,Q]$ *and* $r(u,a)=a$ *for all* $a\in[Q,\overline{a}]$*; and* $r\left(b,\frac{a}{1-\alpha}\right)=Q$ *for all* $a\in\left[0,\frac{Q}{1-\alpha}\right]$ *and* $r\left(b,\frac{a}{1-\alpha}\right)=a$ *for all* $a\in[Q,\overline{a}]$.

And, let $r^* > 0$ satisfy $\Psi(r^*) = \frac{\mu f(\frac{r^*}{(1-\alpha)})}{\mu f(\frac{r^*}{(1-\alpha)})+(1-\mu)f(r^*)} = \frac{\beta+(1-\phi)c}{\chi\phi+c}$. In words, r^* is the smallest ratio r, for which the auditor does not find it optimal to audit, when employers are hiring their first best. Our next result provides an existential result for a marginal equilibrium.

Proposition 2 *A marginal equilibrium exists if and only if there exists an $\widehat{r} \in (0, Q)$ such that the following conditions hold:*

$$\frac{\phi}{1-\phi} = \frac{\Gamma\left(\frac{\widehat{r}}{1-\alpha} - Q\right)}{\Gamma(\widehat{r} - Q)} \text{ for some } \widehat{r} \in [0, Q], \Psi(Q) \leq \frac{\beta + (1-\phi)c}{\chi\phi + c};$$

and any deviation $r' \in [\widehat{r}, Q]$ (out of equilibrium event) is thought to be a discriminator.

Proof. By definition, in any marginal semi-separating equilibrium, there exists a non-empty set of employers such that $r(t, a) < Q$.

Claim 1**: In any semi-separating equilibrium, if** $r(t, a) < Q$, $r(t, a) \in \{a(1-\alpha), a\}$.

Proof of Claim 1 : Suppose $(t, a) = (b, a)$. In this case, he will choose r that satisfies

$$\max\{-\phi P; \Gamma(\alpha a) - (1-\phi)P; \Gamma(z - (1-\alpha)a) - (1-\phi)P\} \text{ for some } z > a.$$

However, the second term is always larger than the third, so we can rewrite this as

$$\max\{-\phi P; \Gamma(\alpha a) - (1-\phi)P\}.$$

A similar argument shows that type a (u, a) employers will hire a, which is the desired result for Claim 1.

Then, there must exists a ratio $\widehat{r} < Q$ (strict inequality follows directly from $\phi < 1$), such that any employer with $a < \widehat{r}$, hires a and for any employer with first best ratio $a \in (\widehat{r}, Q]$, they hire Q. Thus, the employer with $a = \widehat{r}$ must be indifferent between hiring his first best and hiring Q. Further, we know that for any ratio r, there exists a such that $r\left(b, \frac{a}{1-\alpha}\right) = r(u, a) = r$. Therefore, we know at ratio $\widehat{r}$, there is an unbiased employer in state $\widehat{a}$ and a biased employer in state $\frac{\widehat{a}}{1-\alpha}$. To ensure that both employers are indifferent at $\widehat{r}$, the following equations must hold:

$$-(1-\phi)P = \Gamma(Q - \widehat{a})$$

for unbiased types, which implies $P = -\frac{\Gamma(Q-\widehat{a})}{(1-\phi)}$, and

$$-\phi P = \Gamma(Q - \widehat{a}(1-\alpha))$$

for biased types, which implies that $P = -\frac{\Gamma(Q-\widehat{a}(1-\alpha))}{\phi}$. Thus, if $\frac{\phi}{1-\phi} = \frac{\Gamma(Q-\widehat{a}(1-\alpha))}{\Gamma(Q-\widehat{a})}$ both equations are satisfied simultaneously.

To make it optimal for the employer to audit below Q, but not audit at or above, it must be that $\widehat{r} \leq r^*$, $\frac{\mu \int_{\frac{\widehat{r}(Q)}{1-\alpha}}^{\frac{Q}{1-\alpha}} f(a)da}{\mu \int_{\frac{\widehat{r}(Q)}{1-\alpha}}^{\frac{Q}{1-\alpha}} f(a)da + (1-\mu) \int_{\widehat{r}(Q)}^{Q} f(a)da} \leq \frac{\beta+(1-\phi)c}{\chi\phi+c}$, and $Q \geq r^*$, respectively. ■

Proposition 3 *An inclusive equilibrium exists if and only if the following conditions hold:*

$$P > \max\left\{\frac{-\Gamma(-Q)}{1-\phi}, \frac{-\Gamma(-Q)}{\phi}\right\}; \quad \Psi(Q) \leq \frac{\beta + (1-\phi)c}{\chi\phi + c};$$

and any deviation $r' \in [0, Q]$ (out of equilibrium event) is thought to be a discriminator.

Proof. Suppose, by way of contradiction, that there exists separation to the left of Q.

Given Claim 1 above, there must exists a ratio $\widehat{r} < Q$ (strict inequality again follows directly from $\phi < 1$), such that any employer with $a < \widehat{r}$, hires a and for any employer with first best ratio $a \in (\widehat{r}, Q]$, they hire Q. Thus, the employer with $a = \widehat{r}$ must be indifferent between hiring his first best and hiring Q. Further, we know that for any ratio r, there exists a such that $r\left(b, \frac{a}{1-\alpha}\right) = r(u, a) = r$. Therefore, we know at ratio $\widehat{r}$, there is an unbiased employer in state $\widehat{a}$ and a biased employer in state $\frac{\widehat{a}}{1-\alpha}$. To ensure that both employers are indifferent at $\widehat{r}$, the following equations must hold:

$$-(1-\phi) P = \Gamma(Q - \widehat{a})$$

for unbiased types and

$$-\phi P = \Gamma(Q - \widehat{a}(1-\alpha))$$

for biased types. Given $\alpha > 0, P > \max\left\{\frac{\Gamma(a-Q)}{1-\phi}, \frac{\Gamma(a(1-\alpha)-Q)}{\phi}\right\}$, both equations cannot be satisfied simultaneously. Therefore, we need $\widehat{r}_b$ (resp. $\widehat{r}_u$) for biased (resp. unbiased) types, where $\widehat{r}_b < \widehat{r}_u$, that satisfy

$$P = -\frac{\Gamma(Q - \widehat{r}_u)}{(1-\phi)}$$

and

$$P = -\frac{\Gamma(Q - \widehat{r}_b)}{\phi}$$

Thus,

$$\frac{\phi}{1-\phi}=\frac{\Gamma\left(Q-\widehat{r}_b\right)}{\Gamma\left(Q-\widehat{r}_u\right)}$$

which is a contradiction (the r.h.s. is greater than 1 and the l.h.s. is less than 1, by definition). ■

Let

$$g(r)\equiv\frac{\mu f(\frac{r}{(1-\alpha)})}{\mu f(\frac{r}{(1-\alpha)})+(1-\mu)\,f(r)}$$

denote the probability that a profit maximizing ratio r, was hired by a biased employer, where μ denotes the fraction of biased employers in the labor market. I assume that $g^{'}(r)<0$.[10]

Proposition 4 *There exists a vector* $\left(\underline{Q},\ P\right)$ *such that for any* $Q\geq\underline{Q}$*, an inclusive semi-separating equilibrium exists, if all* $r^{'}<Q$ *(out of equilibrium event) are deemed discriminators.*

Proof. Let

$$\underline{Q}\equiv\min\left\{Q:\frac{\mu\int_{Q}^{\frac{Q}{(1-\alpha)}}f(a)da}{\mu\int_{0}^{\frac{Q}{1-\alpha}}f(a)da+(1-\mu)\int_{0}^{Q}f(a)da}\leq\frac{\beta+(1-\phi)\,c}{\chi\phi+c}\right\}$$

It follows directly from Proposition 3 that for any $Q\geq\underline{Q}$, the conditions of the proposition are met, if deviators below Q are thought to be discriminators, and $(1-\phi)\,P$ is sufficiently high. Now, it suffices to show that $\frac{\mu\int_{Q}^{\frac{Q}{1-\alpha}}f(a)da}{\mu\int_{0}^{\frac{Q}{1-\alpha}}f(a)da+(1-\mu)\int_{0}^{Q}f(a)da}$ is decreasing in Q, for $Q\in\left[\max\left\{r^{*},\underline{Q}\right\},\overline{a}\right]$. Since $\Psi\left(Q\right)=\frac{\mu\int_{Q}^{\frac{Q}{1-\alpha}}f(a)da}{\mu\int_{0}^{\frac{Q}{1-\alpha}}f(a)da+(1-\mu)\int_{0}^{Q}f(a)da}$, we have the desired result. ■

Proposition 5 *There exists a vector* $(\underline{Q},P,\widehat{r})$ *such that for any* $Q\geq\underline{Q}$*, a marginal semi-separating equilibrium exists, if all employers who choose* $r^{'}\in[\widehat{r},Q]$ *(out of equilibrium event) are deemed discriminators.*

[10] Taking the first-order derivative, $g^{'}<0$ if and only if:

$$\frac{f^{'}\left(r\right)}{f\left(r\right)}>\frac{f^{'}\left(\frac{r}{1-\alpha}\right)}{\left(1-\alpha\right)f\left(\frac{r}{1-\alpha}\right)}.$$

This is the same condition as $\ln f\left(r\right)$ decreasing in r, which consistent with many distributional assumptions on f.

Proof. Recall, in equilibrium, $\widehat{r}(Q)$. The rest follows directly from Proposition 4. ■

Proposition 6 *All semi-separating equilibria satisfy the intuitive criterion.*

Proof. Using the notation found in Fudenberg and Tirole (ch. 11, p. 448), let $\Theta = \{u, b\} \times [0, \overline{a}]$ denote the set of types, with any particular type denoted $\theta = (t, a)$, $A \in \{0, 1\}$ denote the auditors choice variable, where $A = 1$ if he decides to audit, ξ denote the auditors beliefs, and $u(r^{'}, A, \theta)$ the auditors payoff. Now, define the set of auditor best responses as

$$BR(\Theta, r^{'}) = \bigcup_{\xi:\xi(\Theta|r)=1} BR(\xi, r^{'})$$

where

$$BR(\xi, r^{'}) = \arg\max_{A} \int_{\theta\in\Theta} \xi(\theta \mid r^{'}) u(r^{'}, A, \theta)$$

Let $u_E^*(\theta)$ denote the equilibrium payoff to a type θ employer. Define a set

$$J(r^{'}) = \left\{ \theta : u_E^*(\theta) > \max_{A\in BR(\Theta,r)} u_E(r^{'}, A, \theta) \right\}$$

It is straightforward to see that this set consists of all θ such that $r_\theta^* \geq Q$. Therefore, rewrite $J(r^{'})$ as

$$J(r^{'}) = \{\theta : r_\theta^* \geq Q\}$$

By definition, if for some $r^{'}$, there exists a $\theta^{'} \in \Theta$ such that

$$u_E^*(\theta^{'}) < \min_{A\in BR(\Theta\backslash J(r^{'}),r^{'})} u_E(r^{'}, A, \theta^{'})$$

then the equilibrium fails the intuitive criterion. However, since the set $\Theta \backslash J(r^{'})$ contains discriminators, $BR(\Theta\backslash J(r^{'}), r^{'})$ contains $A = 1$, it follows that

$$\left\{ \theta^{'} \in \Theta : u_E^*(\theta^{'}) < \min_{A\in BR(\Theta\backslash J(r^{'}),r^{'})} u_E(r^{'}, A, \theta^{'}) \right\} = \varnothing$$

which is the desired result. A virtually identical argument proves the analogous result for the set of marginal semi-separating equilibria. *Q.E.D.* ■

APPENDIX B: ADDITIONAL CALCULATIONS

A. Pooling Equilibria

In a pooling equilibrium, all types choose the same action. In particular, a pooling equilibrium exists at r_p if $\forall\ (t,a) \in \{b,u\} \times [0,\overline{a}]$, $r(t,a) = r_p$. In what follows, we prove the existence of a unique pooling equilibrium for our general model.

Proposition 7 *If* $P > \frac{\Gamma(\overline{a})}{(1-\phi)}$, *a unique pooling equilibrium exists at* $r_p = \overline{a}$, *provided that the auditor believes all employers who hire* $r < r_p$ *are discriminating.*

Proof. Suppose $\forall\,(t,a)\ r\,(t,a) = \overline{a}$, and the auditor does not audit any employer with hiring ratio $\overline{a}$, but audits any employer with $r < \overline{a}$. The auditor has no incentive to deviate from this strategy, since he knows that there is no possibility that anyone with $r = \overline{a}$ is discriminating, and it is consistent for him to believe that any deviators are discriminating. For employers, no one has incentive to deviate, since they will be audited and $P > \frac{\Gamma(\overline{a})}{(1-\phi)}$. To establish uniqueness, suppose, by way of contradiction, that there exists a pooling equilibrium at $r^p < \overline{a}$. In this case, $\forall(t,a)$, $r\,(t,a) = r_p$. However, this implies that even an unbiased employer in state $\overline{a}$ does best to deviate. In symbols, this requires that $0 < \Gamma\,(r_p - \overline{a})$, since the auditor will not find it optimal to investigate any employer with $r = \overline{a}$. This contradicts the assumptions on $\Gamma\,(\cdot)$. ■

This proposition provides a knife edge possibility for the existence of pooling equilibrium. The result seems innocuous due to the fact that we do not allow for reverse discrimination in our simple model. This would alleviate such an extreme pooling equilibrium, although we are not certain whether it could guarantee the non-existence of less extreme pooling equilibrium.

B. Non-monotonic Equilibria

In our analysis thus far, we have restricted our attention to monotonic auditing strategies. In this section, we relax that assumption and analyze the existence of equilibria in which the auditor sets non-monotonic threshold strategies.[11] These strategies involve multiple auditing thresholds. This implies that the auditor believes there are certain minority/non-minority hiring ratios that are "just right:" anything too low or too high, is suspect. The final technical result shows that non-monotonic equilibria do not exist in pure strategies.

[11] Technically speaking, these equilibria are in the set of semi-seperating equilibria.

Proposition 8 *No Non-monotonic equilibrium exists.*

Proof. In any non-monotonic equilibrium, there exists at least two auditorial thresholds Q_1 and Q_2, where we assume, without loss of generality, $Q_1 < Q_2$ such that for all $r > Q_2$ the auditor wants to audit and for all $r \leq Q_2$ he does not (else one could assume one auditing threshold without loss).

We know that $Q_2 \in (Q_1, \overline{a}]$. However, it can be shown that $Q_2 = \overline{a}$. Suppose not. Then any employer whose first best hiring ratio $a \geq Q_2$, hires Q_2. However, this can never be optimal for the type $(u, \overline{a})$ employer, given he is guaranteed not to be audited if he hires his first-best. Therefore, the only possible non-monotonic equilibria requires $Q_1 < Q_2 = \overline{a}$. We rule this case out a priori, given this boils down to the auditor using a monotonic strategy (since he must only audit any employer with hiring ratio greater than $\overline{a}$). ■

C. Mixed Strategy Equilibria

For the auditor to play a mixed strategy between multiple thresholds, he must be indifferent between auditing or not at these thresholds. Therefore, $\Psi(r) = \frac{\beta+(1-\phi)c}{\chi\phi+c}$ must have multiple solutions. For tractability, let

$$\Xi \equiv \left\{ r : \Psi(r) = \frac{\beta + (1-\phi)c}{\chi\phi + c} \right\}$$

denote the set of such solutions with its cardinality $\#\Xi \in [1, \infty]$. Finally, let ρ_j denote the probability that an auditor audits at threshold Q_j, where we assume without loss that $Q_1 < Q_2 < ... < Q_{\#\Xi}$

The employers problem is straightforward. A type (t, a) employer hires his first best if that hiring ratio is above $Q_{\#\Xi}$ or, for employers whose profit maximizing hiring ratio is below $Q_{\#\Xi}$, they choose

$$\max \left\{ \begin{array}{c} -(1-\phi)P; \\ \Gamma(Q_1 - a) - (1-\rho_1)(1-\phi)P; \\ \Gamma(Q_2 - a) - (1-\rho_2)(1-\phi)P; \\ \Gamma(Q_{\#\Xi} - a) - \left(1 - \rho_{\#\Xi-1}\right)(1-\phi)P; ...; \\ \Gamma(Q_{\#\Xi} - s) \end{array} \right\}$$

if he is unbiased and

$$\max\{-\phi P; \Gamma(Q_1 - (1-\alpha)a) - (1-\rho_1)\phi P; \Gamma(Q_1 - a) - (1-\rho_1)(1-\phi)P; ...\}$$

for biased employers.

For the auditor, it must be that $\Psi\left(Q_j^r\right) = \frac{\beta+(1-\phi)c}{\chi\phi+c}$, for all $Q_j^r \in \Xi$. Given the general framework for mixed strategy equilibrium, one can check mixed strategies for particular parameter values as needed.