Review of Economic Analysis, forthcoming

Can Public Intervention Mitigate the Welfare Effects of Discrimination?

A Reduced-Form Surplus-Sharing Search Model

Félix Dopchie*

June, 2026

Abstract

Despite strong legal frameworks prohibiting discrimination, unequal treatment across protected and unprotected groups remains widespread in labor markets. This paper studies whether anti-discrimination policies improve welfare when discriminatory practices are flexible across margins. I develop a reduced-form search model with endogenous vacancy creation and taste-based disutility, in which workers face a wage–probability trade-off. I analyze two stylized interventions: an Equal Pay Policy, which constrains wages, and a Quota Policy, which constrains job-finding probabilities. The paper shows that wage discrimination and hiring-probability discrimination are jointly determined and can substitute for one another. As a result, policies may achieve their formal equality objectives while reducing welfare. Equal Pay can eliminate wage differences but worsen protected workers' job-finding probabilities when vacancy creation is not sufficiently elastic, whereas a quota can equalize hiring probabilities and indirectly equalize wages while pooling workers into a lower common allocation. The two policies also differ in incidence: Equal Pay pools wages but leaves employment probabilities group-specific, while the quota pools the full wage–probability allocation. The main policy implication is that anti-discrimination policies should be evaluated as equilibrium interventions affecting wages, vacancies, and employment probabilities, not only by their direct effect on the targeted margin.

Keywords: Search frictions, discrimination, surplus sharing, equal pay, quotas, welfare

JEL Classifications: I38, J71, J78

1 Introduction

Despite extensive legal frameworks aiming to reduce disparities in labor markets (e.g., European Union (2000b), which explicitly prohibits discrimination in hiring, promotion, and remuneration), discrimination remains a persistent feature of labor-market outcomes.

*UCLouvain - Economics School of Louvain (ESL), Belgium: felixdopchie@hotmail.com

Empirical studies continue to document substantial gaps across demographic groups, including race, gender, ethnicity, and origin (e.g., Bertrand and Mullainathan (2004) and Blau and Kahn (2017)). This persistence raises a central policy question: *To what extent can public intervention mitigate the welfare effects of labor-market discrimination when discriminatory practices are flexible and substitutable?* This paper addresses this question through a theoretical lens. It studies whether standard anti-discrimination policies, such as Equal Pay and Quota Policies, improve the welfare of protected groups when firms and workers can re-optimize across different margins of discrimination.

A rich body of literature has documented the various forms that labor-market discrimination can take. Initially, the focus was primarily on wage disparities, but more recent work has highlighted more complex and less observable mechanisms. For example, Blau and Kahn (2017) argue that gender-based discrimination has evolved from overt wage gaps to more subtle forms such as differences in job quality. Similarly, Bertrand and Mullainathan (2004) use a natural experiment based on *résumés* to provide compelling evidence of racial discrimination in hiring decisions. They document that applicants with African-American names receive about 50% fewer callbacks in the United States. Recent studies (e.g., Bartos et al. (2013), Bertrand and Duflo (2017), Galos and Coppock (2023)) examine the interaction between different forms of discrimination, showing that their effects are often cumulative rather than isolated. These contributions build on the foundational work of Becker (1957), who introduced the concept of taste-based discrimination, and they highlight that discrimination is deeply embedded in labor-market dynamics. Despite these advances, identifying the magnitude of discrimination remains empirically challenging. As Charles and Guryan (2011) note, traditional decomposition methods (e.g., Oaxaca (1973), Gelbach (2016)) often face limitations in isolating the discriminatory component from other unobserved factors. Moreover, while natural experiments offer cleaner identification, they may rely on artificial setups that limit external validity. Neumark (2018), for instance, critiques the reliance on *résumé*-based field experiments, arguing that they often fail to capture the full complexity of hiring behavior in real-world firms.

Although still relatively underexplored (Bertrand and Duflo (2017)), a growing body of research moves beyond detection and measurement to assess the welfare effects of discrimination and the role of public policy. Notably, Donohue and Heckman (1991), Heckman and Payner (1989), and Heckman and Siegelman (1993) show that discrimination can generate inefficiencies by distorting the allocation of talent and reducing aggregate productivity. In this view, labor markets that systematically disadvantage certain groups not only create inequities but also misallocate resources. Public intervention, when carefully designed, may therefore correct such inefficiencies, especially when combined with appropriate market mechanisms.

Within this broader conversation, affirmative action and Equal Pay Policies have received particular attention. For example, Holzer and Neumark (2000) assess the general equilibrium impacts of affirmative action and highlight several trade-offs. While such policies can improve representation and reduce wage disparities, they may also distort wage formation or hiring incentives. These insights emphasize that the effectiveness of anti-discrimination policy depends not only on the legal objective of the intervention, but also on the underlying structure of discrimination and on the behavioral responses of firms and workers.

Despite these important contributions, a gap remains in the theoretical literature: most models do not explicitly consider the *substitutability* of discriminatory practices. In practice, discrimination need not be tied to a single fixed margin. A firm that is prevented from discriminating in wages may adjust by discriminating in hiring, by altering job quality, or by changing the types of vacancies it creates. This paper aims to fill this gap by studying how policy effectiveness changes when discriminatory mechanisms are flexible and strategically substitutable. The central message is that a policy cannot be evaluated solely by asking whether it eliminates the margin it directly targets. It must also be evaluated by asking how the equilibrium reallocates discrimination, employment probabilities, wages, and welfare across groups.

To explore this question, I develop a reduced-form search environment with submarket choice, endogenous vacancy creation, and taste-based disutility in the spirit of Becker (1957). Workers choose across submarkets that differ in wages and job-finding probabilities, while firms create vacancies taking the wage rule as given. Wage determination is represented by a tightness-dependent surplus-sharing rule. This rule is imposed as a tractable reduced-form device and is not derived as the outcome of a full wage-posting, bargaining, or competitive-search game. The model therefore focuses on how anti-discrimination policies reshape the wage–probability allocation, rather than on the microfoundation of wage posting itself.

The primary public interventions examined in this paper are stylized versions of an Equal Pay Policy, akin to those analyzed by Council et al. (1981) and Polachek (2019), and a Quota Policy. In practice, such policies are imposed at the firm or organization level. The model does not explicitly describe within-firm wage setting or firm-level legal compliance. Instead, the Equal Pay exercise is modeled as a reduced-form restriction on group-based wage differences, while the Quota exercise is modeled as a reduced-form restriction on group-based differences in job-finding probabilities. These policy constraints should therefore be interpreted as tractable aggregate representations of firm-level policies, not as literal descriptions of legal implementation within firms. The distinction between the two interventions is central for policy design, because the model shows that regulating one margin can change equilibrium behavior on the other.

Within this reduced-form environment, the paper delivers three policy-relevant insights. First, wage differences and hiring-probability differences are jointly determined. A policy that restricts one margin may therefore induce adjustment on the other margin. This is important for policy evaluation because discrimination should not be treated as a one-dimensional distortion. A reform that eliminates an observed wage gap may still worsen employment access, while a reform that equalizes hiring probabilities may affect wages through vacancy creation and surplus sharing.

Second, Equal Pay can reduce expected welfare under incomplete vacancy adjustment. In the model, wage equalization can make the disutility submarket unattractive or unprofitable. The disappearance of this submarket reduces the set of active vacancies available in the economy. Non-disutility firms may create additional vacancies in response, but because vacancy creation is convex, this response is not infinitely elastic and need not fully offset the lost vacancies. Equal Pay therefore generates both a wage effect and a congestion effect. The wage effect raises the protected workers' wage relative to some laissez-faire allocations, while wage pooling may reduce the wage component received by unprotected workers. The congestion effect operates through equilibrium job-finding probabilities and can lower expected employment rents for one or both groups. Thus, the policy cost is not imposed only on protected workers; it can also be borne by unprotected workers through lower job-finding probabilities or a lower wage component.

Third, Equal Pay and Quota Policies have different incidence structures. Equal Pay is a wage-pooling policy: it imposes a common wage while allowing job-finding probabilities to remain group-specific. The Quota Policy is an allocation-pooling policy: it restricts job-finding-probability differences and, in the closed-form version of the model, indirectly generates wage equality across active non-disutility submarkets. These policies may therefore achieve formal equality objectives while having different effects on expected welfare and on the distribution of policy costs across groups.

The central mechanism is therefore not that equality is intrinsically costly. Rather, anti-discrimination policies may backfire when they constrain one discriminatory margin while leaving firms and workers free to re-optimize along other margins. In that sense, the paper's main contribution is to show that policy evaluation must account for the substitutability of discriminatory practices, not only for the direct effect of the policy on the targeted margin. A narrow evaluation that focuses only on wages, or only on hiring probabilities, may miss the equilibrium response that determines whether the policy actually improves welfare.

Equal Pay illustrates this mechanism most directly. By construction, it restricts wage differences across otherwise identical workers. In the reduced-form environment, this can make the disutility submarket unattractive or unprofitable for protected workers. Protected workers may then redirect their applications toward the non-disutility submarket, while non-disutility firms adjust vacancy creation. Because vacancy creation is convex, the expansion of non-disutility vacancies need not fully compensate for the disappearance of disutility vacancies. Equal Pay therefore has two opposing effects: it affects wages through wage pooling, and it affects expected employment rents through equilibrium job-finding probabilities. Whether the policy improves or worsens welfare depends on which effect dominates. This mechanism corresponds to an Equal Pay experiment that constrains the wage margin without mechanically integrating protected and unprotected workers into a single common search market. In the constrained equilibrium, wages, vacancies, tightness, and job-finding probabilities are recalculated. The adverse probability effect arises when convex vacancy creation prevents non-disutility firms from fully replacing the search opportunities lost when the disutility submarket becomes inactive.

These findings are relevant for the design of anti-discrimination policy. They do not imply that public intervention is undesirable. Rather, they show that policies targeting a single discriminatory margin may be insufficient when firms can substitute across margins. Effective policy design must therefore consider the full set of behavioral responses generated by regulation. In particular, wage regulation may need to be complemented by policies affecting vacancy creation, enforcement against hiring discrimination, or instruments that directly target the employment margin. More broadly, the analysis suggests that anti-discrimination policy should be evaluated not only by whether it removes an observed gap, but also by whether it improves the expected welfare of the groups it aims to protect.

Finally, interventions such as Equal Pay and Quota Policies can reduce aggregate welfare, defined in an *ex-ante utilitarian welfarist* sense, that is, as the sum of expected utilities across the population. This welfare reduction arises from disrupting the equilibrium allocation that emerges under *laissez-faire*, an allocation jointly determined by both wage and hiring discrimination. These findings are consistent with, and extend, the work of Heckman and Siegelman (1993) and Lang, Manove, and Dickens (2005), by showing that the effectiveness of anti-discrimination policies crucially depends on the interaction between different forms of discrimination.

The structure of the paper is as follows. Section 2 reviews the theoretical foundations of discrimination in labor markets and introduces the modeling framework. Section 3 presents the formal model and characterizes the equilibrium. Section 4 provides a closed-form characterization of the equilibrium and derives transparent welfare conditions for Equal Pay and Quota Policies. Section 5 studies aggregate welfare and policy incidence across groups. Section 6 discusses the implications for public policy and concludes.

2 Search Models and Discrimination

This section reviews the theoretical literature on labor-market discrimination, explains why search frictions are a natural framework for studying its persistence, and positions the modeling strategy adopted in this paper. I first clarify the definition of discrimination used throughout the analysis. I then discuss how search models depart from the perfectly competitive benchmark and why this departure is important for understanding persistent discrimination. Finally, I relate the paper to the literature on public intervention and welfare in discriminatory labor markets.

Discrimination is not defined uniformly across disciplines. In sociology, unequal outcomes are often viewed as evidence of discrimination when they reflect or reinforce disadvantages associated with race, origin, gender, or other protected characteristics (Reskin (2000), Pager and Shepherd (2008)). Under this broader view, discrimination may operate even when group identity is not explicitly used as a decision criterion, since pre-existing inequalities can generate cumulative disadvantages. In contrast, economics typically adopts a narrower definition: discrimination occurs when equally productive workers are treated differently because of non-economic characteristics such as race, gender, ethnicity, or origin (Becker (1957), Phelps (1972))[1]. While the European Union adopts a broad legal definition of discrimination (European Union (2000b), European Union (2000a)), this paper follows the economic approach. Discrimination is defined as unequal treatment of equally productive workers on the basis of group identity. Under perfect competition, such unequal treatment is inefficient, since hiring and wage decisions should depend on productivity rather than group identity (Arrow (1972), Becker (1957), Kaas and Lu (2010)).

Starting with Becker (1957), a large theoretical literature has studied the origins and consequences of labor-market discrimination. Becker's taste-based model is the natural benchmark for this paper. In that framework, prejudiced employers incur a disutility from hiring workers from a protected group. However, under perfect competition and in the absence of market frictions, disutility firms are eventually driven out of the market because non-disutility firms can hire equally productive protected workers at lower effective cost (Becker (1957), Arrow (1972)). As a result, discrimination should disappear in the long-run competitive equilibrium. Subsequent contributions have extended the analysis to statistical discrimination and belief-based mechanisms (Arrow (1971), Phelps (1972)), bounded rationality and biased updating (Bohren, Imas, and Rosenberg (2019)), stereotype-based beliefs (Bordalo (2016)), and network-based or homophily-driven mechanisms (Calvo-Armengol and Jackson (2004)).

Search theory provides a natural way to depart from the frictionless competitive benchmark. In search models, workers and firms do not meet instantaneously, and labor-market outcomes depend on matching frictions, application behavior, vacancy creation, and market tightness. These frictions help explain deviations from the Walrasian benchmark (Arrow and Debreu (1954); originally formulated by Walras (1877)) and account for features such as unemployment, firm market power, wage dispersion, and mismatch (Rogerson, Shimer, and Wright (2005)). In the context of discrimination, frictions are especially important because they can prevent the Beckerian mechanism from fully eliminating disutility firms. If non-disutility firms cannot expand employment without limit, disutility firms may remain active in equilibrium despite being less efficient.

Building on foundational random search models (Diamond (1971) and Diamond (1982)), the literature has developed several frameworks to study unemployment, wage dispersion, vacancy creation, and matching frictions, including on-the-job search (Pissarides (2000)), competitive search (Moen (1997)), and job creation and destruction (Mortensen

[1]For a comprehensive comparison, see Pager (2008).

and Pissarides (1994)). Burdett and Mortensen (1998) is best understood here as a canonical search model of wage dispersion rather than as a model of discrimination. Search-theoretic models of discrimination include, among others, Black (1995), Bowlus and Eckstein (2002), Flabbi (2010), and Lang, Manove, and Dickens (2005). These papers show how search frictions, bargaining, wage posting, or employer prejudice can generate persistent group differences in wages, employment, or welfare.

This paper builds on this search-friction tradition by introducing taste-based disutility into a tractable submarket-choice environment. Firms differ in whether they incur a disutility cost from hiring protected workers. Protected workers can choose across submarkets associated with different wage–probability profiles, while firms create vacancies subject to convex vacancy-creation costs. The model does not solve a full wage-posting, bargaining, or competitive-search equilibrium. Instead, wages are determined by a tightness-dependent surplus-sharing rule. This reduced-form structure preserves the wage–probability trade-off central to search environments while keeping the policy analysis analytically tractable.

The advantage of this approach is that it makes it possible to study how public intervention affects the full wage–probability allocation in a discriminatory labor market. This is important because anti-discrimination policies typically target one observable margin, such as wages or hiring rates, while firms may adjust along other margins. Traditional models often attach discrimination to a single instrument or outcome (Becker (1957), Kaas and Lu (2010), Rosen (1974)). By contrast, the framework developed here allows wage discrimination and hiring discrimination to be jointly determined and potentially substitutable. Restricting one margin of discrimination changes the value of the remaining margins and therefore changes equilibrium application and vacancy-posting behavior.

Welfare analyses of discrimination and anti-discrimination policy remain relatively scarce, as emphasized by Bertrand and Duflo (2017). On the empirical side, studies such as Neumark (2018) highlight the importance of legal frameworks against hiring discrimination. Donohue (2005) argues that anti-discrimination policies can improve welfare for discriminated groups in both taste-based and statistical discrimination settings, although political support for such policies is often related to egalitarian preferences (Ziller and Helbling (2019)). On the theoretical side, Kaas and Lu (2010) show that Equal Pay Policy can increase inequality among discriminated workers in a random search model, highlighting the possibility that well-intended interventions may have adverse distributional effects. Similarly, Bowlus and Eckstein (2002) and other search-theoretic models of discrimination show that restrictions on one form of discrimination may induce firms to adjust along alternative margins, potentially affecting market efficiency.

This paper contributes to this literature by studying how wage and hiring-probability margins interact when anti-discrimination policy restricts one of them. The mechanism differs from models in which discrimination is purely wage-based, purely hiring-based, or eliminated once a single policy is imposed. In the present reduced-form environment, firms and workers interact in submarkets characterized by different wage–probability combinations. A policy may therefore be successful in a narrow legal sense while failing in welfare terms. Equal Pay can eliminate wage differences but worsen hiring prospects for protected workers when vacancy creation is insufficiently elastic. A quota can eliminate hiring-probability differences and, in the closed-form model, indirectly eliminate wage differences, but it may pool workers into a lower common allocation.

The question of efficiency in discriminatory labor markets has also been addressed by Lang, Manove, and Dickens (2005) and Lang and Lehmann (2012), who develop search models with disutility parameters close to Becker's approach and show that reducing discrimination can increase total productivity and minority employment. Their models emphasize the productivity and employment gains that may arise from combating discrimination. However, their reliance on bargaining structures makes some welfare comparisons less transparent. The reduced-form surplus-sharing approach used in this paper avoids modeling the bargaining or wage-posting game explicitly. Instead, it characterizes wages, vacancy creation, and job-finding probabilities through submarket-level conditions. This makes the welfare effects of Equal Pay and Quota Policies transparent, while also limiting the scope of the analysis: the surplus-sharing rule should be read as an assumption, not as an equilibrium outcome of a fully specified wage-setting game. This allows the welfare effects of Equal Pay and Quota Policies to be expressed directly in terms of the wage–probability trade-off faced by workers.

Finally, the model clarifies why discrimination can persist despite the Beckerian prediction that competition should eliminate it in the long run. The standard Becker model relies on perfect competition and frictionless adjustment. In contrast, the model developed in Section 3 combines taste-based disutility with matching frictions and convex vacancy-creation costs. First, some firms have a firm-specific disutility from hiring protected workers, capturing prejudice that is independent of productivity. Second, matching frictions and reduced-form segmentation generate distinct submarkets with different wage–probability allocations. Third, because vacancy creation is costly and not infinitely elastic, non-disutility firms cannot absorb all protected workers at no cost. Given the constant-returns matching technology (Petrongolo and Pissarides (2001)) and the equilibrium conditions described in Section 3.1, outcomes can differ across submarkets, preventing competitive forces from fully eradicating discrimination. This persistence is essential for the policy analysis that follows: if discrimination vanished automatically, Equal Pay and Quota Policies would have no meaningful role to play.

3 The Model

This section develops a static reduced-form search model with submarket choice, matching frictions, endogenous vacancy creation, and taste-based disutility. Workers are divided into unprotected and protected workers. Unprotected workers search in the unprotected-worker submarket. Protected workers can choose between a submarket associated with non-disutility firms and a submarket associated with disutility firms. Discrimination can persist because vacancy creation is endogenous but not infinitely elastic: the marginal cost of creating vacancies is increasing. This prevents non-disutility vacancies from becoming infinitely abundant and therefore avoids the frictionless Beckerian outcome in which disutility firms are eliminated by unrestricted entry.

The environment should not be interpreted as a full competitive-search or directed-search wage-posting equilibrium. Wages are not posted firm by firm. Instead, they are determined by a tightness-dependent surplus-sharing rule that summarizes wage determination within each active submarket. This rule is imposed as a tractable reduced-form assumption; it is not derived as the equilibrium outcome of a full wage-posting, bargaining, or competitive-search game. The purpose of the model is therefore to study how public interventions reshape equilibrium wage–probability allocations, not to microfound wage posting itself.

3.1 Environment

Workers. Let I denote the set of workers, endowed with a σ-algebra $\mathcal{I}$. Let λ be a finite measure on $(I, \mathcal{I})$. Workers are partitioned into unprotected and protected workers: $I = I_U \cup I_P$, with $I_U \cap I_P = \varnothing$. The total mass of workers is normalized to one. The mass of unprotected workers is $\lambda(I_U) = \mu$, with $\mu \in (0, 1)$, and the mass of protected workers is $\lambda(I_P) = 1 - \mu$.

All workers have productivity $y > 0$ when employed and receive an outside option $b > 0$ when unemployed. I assume $y > b$. Workers have identical preferences over wages. These preferences admit a Bernoulli utility representation $u : \mathbb{R}_+ \to \mathbb{R}$, with $u'(w) > 0$ and $u''(w) \leq 0$. Applying to a vacancy entails a cost $\delta \geq 0$, and each worker can apply at most once.

There are three submarkets. Unprotected workers apply to the unprotected-worker submarket, denoted by N. Protected workers can apply either to the protected-worker submarket associated with non-disutility firms, denoted by 0, or to the protected-worker submarket associated with disutility firms, denoted by d. Protected workers are not assumed to observe firms' disutility parameters directly. The reduced-form assumption is instead that they can distinguish application channels associated with different wage–probability profiles. These channels may reflect occupations, locations, recruitment networks, referral structures, or observed histories of treatment. Let $\mathcal{S}_U = \{N\}$, $\mathcal{S}_P = \{0, d\}$, and $\mathcal{S} = \{N, 0, d\}$.

The restriction that protected and unprotected workers search over different sets of submarkets is a reduced-form segmentation assumption. It is not derived endogenously. The assumption captures institutional, informational, occupational, spatial, or network barriers that prevent full arbitrage across all vacancies. It should not be interpreted as saying that protected workers are less productive or intrinsically unable to perform the same jobs as unprotected workers. Rather, it isolates the mechanism studied in the paper: when protected workers face different effective search opportunities, policy constraints on wages or hiring probabilities can induce adjustment along other margins. Without such segmentation, a richer model in which all firms can hire all worker types and potentially offer type-specific wages would be required.

If a worker applies to submarket $s \in \mathcal{S}$, receives wage w_s when employed, and faces job-finding probability p_s, expected utility is

$$U_s = p_s u(w_s) + (1 - p_s)u(b) - \delta. \tag{1}$$

Not applying gives utility $u(b)$. Workers therefore apply only if the best available submarket delivers expected utility weakly above $u(b)$. In the main analysis, I focus on interior equilibria in which all workers apply. This restriction keeps attention on the wage–probability trade-off generated by discrimination and public intervention. If one group did not apply, the welfare comparison would collapse into a participation-margin result rather than a comparison of equilibrium allocations across active submarkets.

Let $a_i \in \mathcal{S}_j \cup \{\varnothing\}$ denote the application choice of worker $i \in I_j$, where $j \in \{U, P\}$ and $\varnothing$ denotes non-application. The mass of applications in submarket s is

$$A_s = \lambda\left(\{i \in I : a_i = s\}\right).$$

In an interior equilibrium,

$$A_N = \mu, \qquad A_0 + A_d = 1 - \mu. \tag{2}$$

Firms. Let K denote the set of firms, endowed with a σ-algebra $\mathcal{K}$. Let ζ be a finite measure on $(K, \mathcal{K})$. Firms are partitioned into non-disutility and disutility firms: $K = K_0 \cup K_d$, with $K_0 \cap K_d = \varnothing$. Firms in K_0 do not incur a taste-based disutility from employing protected workers. Firms in K_d incur a disutility cost $d > 0$ when employing a protected worker. The term non-disutility firm therefore means only that the firm has no taste-based disutility parameter. It does not imply that the firm necessarily offers the same wage–probability allocation to protected and unprotected workers. In particular, if protected workers have a weaker effective outside option or face less favorable search conditions, non-disutility firms may still offer them less favorable terms. Let c_s denote the firm-side psychological cost associated with a filled vacancy in submarket s. Thus $c_N = 0$, $c_0 = 0$, and $c_d = d$.

Consider first an individual firm posting v_k^s vacancies in submarket s. Taking the vacancy-filling probability q_s and the wage w_s as given, firm k's expected payoff from these vacancies is

$$\pi_k^s = q_s v_k^s (y - c_s - w_s) - C(v_k^s). \tag{3}$$

The term $y - c_s - w_s$ is the firm's net payoff from a filled vacancy before vacancy-creation costs. The term $C(v_k^s)$ is the cost of posting v_k^s vacancies in submarket s.

Aggregating across firms active in submarket s, let V_s denote the total mass of vacancies posted in that submarket. The aggregate net payoff from vacancies in submarket s is

$$\Pi_s = q_s V_s (y - c_s - w_s) - C(V_s). \tag{4}$$

The aggregate vacancy cost satisfies $C(0) = 0$, $C'(V_s) > 0$, and $C''(V_s) > 0$. The strict convexity of C captures limited managerial capacity, limited entrepreneurial opportunities, or adjustment costs in vacancy creation. Thus vacancy creation is endogenous, but not infinitely elastic.

Competitive vacancy creation is characterized by

$$V_s \geq 0, \qquad C'(V_s) - q_s(y - c_s - w_s) \geq 0, \qquad V_s \left[C'(V_s) - q_s(y - c_s - w_s)\right] = 0. \tag{5}$$

If $V_s > 0$, this reduces to $C'(V_s) = q_s(y - c_s - w_s)$.

Given the outside option b, the total match surplus in submarket s is

$$S_s = y - b - c_s.$$

Thus $S_N = S_0 = y - b$ and $S_d = y - b - d$. I assume throughout that $S_d > 0$, so that disutility firms can generate positive surplus when employing protected workers. This surplus expression is used below to determine wages under the surplus-sharing rule.

Matching and surplus sharing. In each submarket $s \in \mathcal{S}$, matches are produced by a matching function $m_s = \mathcal{M}(A_s, V_s)$. The function $\mathcal{M} : \mathbb{R}_+^2 \to \mathbb{R}_+$ is continuous, increasing in both arguments, concave, and homogeneous of degree one.

The job-finding probability is

$$p_s = \frac{\mathcal{M}(A_s, V_s)}{A_s}, \tag{6}$$

and the vacancy-filling probability is

$$q_s = \frac{\mathcal{M}(A_s, V_s)}{V_s}. \tag{7}$$

Let $\theta_s = V_s/A_s$ denote market tightness in submarket $s \in \mathcal{S}$. By constant returns to scale, $p_s = p(\theta_s)$ and $q_s = q(\theta_s) = p(\theta_s)/\theta_s$. I assume $p'(\theta) > 0$ and $q'(\theta) < 0$. A tighter submarket raises workers' job-finding probability and lowers firms' vacancy-filling probability.

The share of surplus accruing to workers is allowed to depend on market tightness. Let $\beta : \mathbb{R}_+ \to (0, 1)$ be continuously differentiable, with $\beta'(\theta) > 0$. Hence, workers capture a larger share of the match surplus in tighter submarkets, where vacancies are relatively abundant and workers have stronger outside opportunities. Conversely, when applications are abundant relative to vacancies, firms have a stronger market position and workers capture a lower share of the surplus. The wage in submarket s is therefore determined by

$$w_s = b + \beta(\theta_s)S_s. \tag{8}$$

This formulation nests the fixed-sharing case as a special case, while allowing policy-induced changes in market tightness to affect both job-finding probabilities and wages.

3.2 *Laissez-faire* Equilibrium

Under *laissez-faire*, wages are determined separately in each submarket. Using (8), the wage in submarket s is

$$w_s^{LF} = b + \beta(\theta_s^{LF})S_s, \qquad s \in \mathcal{S}. \tag{9}$$

Since $S_d < S_0$, a match between a protected worker and a disutility firm generates less surplus than a match between a protected worker and a non-disutility firm. However, because the worker's share $\beta(\theta_s)$ also depends on market tightness, wage discrimination under *laissez-faire* requires

$$\beta(\theta_0^{LF})S_0 > \beta(\theta_d^{LF})S_d.$$

This condition says that the surplus loss caused by the disutility cost d dominates any compensating increase in the worker's surplus share in the disutility submarket. When it holds, $w_0^{LF} > w_d^{LF}$.

Definition 1. *A laissez-faire equilibrium is a collection*

$$\left\{A_s^{LF}, V_s^{LF}, p_s^{LF}, q_s^{LF}, w_s^{LF}\right\}_{s\in\mathcal{S}}$$

such that:

1. *Wages satisfy* (9).
2. *Matching probabilities satisfy* (6) *and* (7).
3. *Workers choose submarkets optimally. In particular, if both protected-worker submarkets are active, then*

$$p_0^{LF}\left[u(w_0^{LF}) - u(b)\right] = p_d^{LF}\left[u(w_d^{LF}) - u(b)\right]. \tag{10}$$

4. *Applications clear as in* (2).
5. *Vacancy creation satisfies* (5).

The *laissez-faire* equilibrium considered in the rest of the paper is an interior equilibrium in which the unprotected-worker submarket and both protected-worker submarkets are active. This means $A_N^{LF} > 0$, $A_0^{LF} > 0$, $A_d^{LF} > 0$, and $V_N^{LF} > 0$, $V_0^{LF} > 0$, $V_d^{LF} > 0$. This is the relevant case for studying persistent discrimination, because disutility firms remain present in the protected-worker labor market and protected workers are willing to apply to them.

If both protected-worker submarkets are active, protected-worker expected utility under *laissez-faire* is

$$U_P^{LF} = u(b) + p_0^{LF} \left[u(w_0^{LF}) - u(b)\right] - \delta, \tag{11}$$

where the same value is obtained using the disutility submarket by (10). Unprotected-worker expected utility is

$$U_U^{LF} = u(b) + p_N^{LF} \left[u(w_N^{LF}) - u(b)\right] - \delta. \tag{12}$$

Proposition 1. *Suppose both protected-worker submarkets are active under laissez-faire and*

$$\beta(\theta_0^{LF})S_0 > \beta(\theta_d^{LF})S_d.$$

Then the equilibrium features wage discrimination and compensating differences in job-finding probabilities:

$$w_0^{LF} > w_d^{LF}, \qquad p_0^{LF} < p_d^{LF}.$$

Proof. See Appendix A.1 □

Comparison with the unprotected-worker submarket. The *laissez-faire* allocation also allows protected workers employed by non-disutility firms to receive a different wage from unprotected workers. The relevant comparison is between

$$w_N^{LF} = b + \beta(\theta_N^{LF})S_0,$$

and

$$w_0^{LF} = b + \beta(\theta_0^{LF})S_0.$$

Since both active non-disutility submarkets generate the same match surplus $S_0 = y - b$, any wage difference comes from differences in market tightness and therefore from differences in effective search opportunities. If $\theta_0^{LF} < \theta_N^{LF}$, then $w_0^{LF} < w_N^{LF}$. Thus, even firms without a taste-based disutility may offer protected workers a lower wage when protected workers have a weaker effective outside option or face less favorable search conditions. This is why the model distinguishes between disutility and non-disutility firms, rather than between discriminatory and non-discriminatory firms.

The persistence of discrimination requires the disutility submarket to be active. This requires both $A_d^{LF} > 0$ and $V_d^{LF} > 0$. The first condition is the worker-side participation condition for the disutility submarket. The second is the firm-side vacancy-creation condition. From (5), a necessary condition for $V_d^{LF} > 0$ is

$$q_d^{LF}(y - d - w_d^{LF}) > C'(0). \tag{13}$$

Using (9), this can be written as

$$q_d^{LF}\left[1-\beta(\theta_d^{LF})\right]S_d > C'(0).$$

Thus, $S_d > 0$ is necessary but not sufficient. Disutility firms survive in the long run only if they can both attract protected-worker applications and cover the marginal cost of vacancy creation. The strict convexity of vacancy costs is important here: it prevents non-disutility firms from expanding vacancies without limit and absorbing all protected workers.

3.3 Equal Pay Policy Equilibrium

The Equal Pay exercise is a stylized aggregate representation of a firm-level equal-pay constraint. In practice, equal-pay rules restrict wage differences within firms among otherwise comparable workers. The model does not explicitly describe within-firm wage setting. Instead, it captures the reduced-form effect of restricting group-based wage differences by imposing a common wage across the relevant active submarkets. Since all workers have productivity y, the policy constraint is written as

$$w_N^{EP} = w_0^{EP} = w_d^{EP} \equiv w^{EP}. \tag{14}$$

The common wage is not fixed at its *laissez-faire* unprotected-worker value. It is determined in the constrained equilibrium. Equation (15) should be read as an aggregate reduced-form wage rule: the total wage surplus paid to workers is consistent with the tightness-dependent surplus shares generated across active matches. The common wage is therefore not required to satisfy the unconstrained submarket-specific wage rule separately in each active submarket. This distinction is important. If the unconstrained wage rule had to hold separately in each active non-disutility submarket, wage equality would mechanically imply equal tightness and hence equal job-finding probabilities. The Equal Pay exercise considered here instead asks what happens when the wage margin is constrained while job-finding probabilities remain free to adjust.

$$\left(w^{EP}-b\right)\sum_{s\in\mathcal{S}}\mathcal{M}(A_s^{EP},V_s^{EP}) = \sum_{s\in\mathcal{S}}\beta(\theta_s^{EP})\mathcal{M}(A_s^{EP},V_s^{EP})S_s. \tag{15}$$

Equivalently, w^{EP} is the match-weighted common wage consistent with Equal Pay and the endogenous surplus-sharing rule. Under *laissez-faire*, each active submarket satisfies (8). Under Equal Pay, by contrast, the policy imposes a common wage across submarkets. The common wage is not required to satisfy the unconstrained wage rule separately in each active submarket. If it were, then, since the active unprotected-worker and non-disutility protected-worker submarkets have the same surplus $S_N = S_0$, wage equality would imply $\beta(\theta_N^{EP}) = \beta(\theta_0^{EP})$, and hence $\theta_N^{EP} = \theta_0^{EP}$ whenever β is strictly increasing. This would mechanically imply equal job-finding probabilities. The Equal Pay equilibrium considered here instead requires the common wage to satisfy the aggregate surplus-sharing condition in (15). Therefore, wage equality can coexist with different market tightness levels and different job-finding probabilities across the active unprotected-worker and protected-worker submarkets.

The following proposition characterizes the relevant Equal Pay equilibrium under the reduced-form wage constraint. I impose a tie-breaking convention: whenever protected workers are indifferent between the disutility and non-disutility protected-worker submarkets, they apply to the non-disutility submarket. This convention selects the equilibrium in which the disutility submarket is not sustained by worker indifference alone.

Proposition 2. *Suppose Equal Pay imposes a common wage w^{EP}, protected workers break ties in favor of the non-disutility protected-worker submarket, and*

$$q(\theta_0^{EP})(y - d - w^{EP}) \leq C'(0). \tag{16}$$

Then, in the relevant Equal Pay equilibrium, the disutility protected-worker submarket is inactive:

$$A_d^{EP} = 0, \qquad V_d^{EP} = 0. \tag{17}$$

Proof. See Appendix A.2 □

Condition (16) is a no-profitable-entry condition for the disutility submarket under Equal Pay. It states that, at the tightness required to compete with the non-disutility protected-worker submarket, the expected return from the first disutility vacancy is no larger than its marginal cost. Thus, disutility firms cannot profitably compensate their lower surplus by offering a sufficiently high job-finding probability at the common Equal Pay wage. Protected workers therefore apply to the non-disutility submarket, so that

$$A_0^{EP} = 1 - \mu. \tag{18}$$

Non-disutility firms may respond by creating additional vacancies, but vacancy creation is convex. The additional vacancies need not fully offset the congestion created by the reallocation of protected workers toward submarket 0.

Definition 2. *An Equal Pay equilibrium is a collection*

$$\left\{A_s^{EP}, V_s^{EP}, p_s^{EP}, q_s^{EP}, w^{EP}\right\}_{s \in \mathcal{S}}$$

such that:

1. *Wages satisfy the Equal Pay constraint* (14) *and the constrained wage rule* (15).
2. *The disutility protected-worker submarket is inactive, as in* (17).
3. *Matching probabilities satisfy* (6) *and* (7) *on active submarkets.*
4. *Applications clear, with $A_N^{EP} = \mu$, $A_0^{EP} = 1 - \mu$, and $A_d^{EP} = 0$.*
5. *Vacancy creation satisfies* (5) *on active submarkets, with $w_s = w^{EP}$.*

Protected-worker expected utility under Equal Pay is

$$U_P^{EP} = u(b) + p_0^{EP}\left[u(w^{EP}) - u(b)\right] - \delta. \tag{19}$$

Unprotected-worker expected utility under Equal Pay is

$$U_U^{EP} = u(b) + p_N^{EP}\left[u(w^{EP}) - u(b)\right] - \delta. \tag{20}$$

Proposition 3. *Equal Pay reduces protected-worker welfare relative to laissez-faire if and only if*

$$p_0^{EP}\left[u(w^{EP}) - u(b)\right] < p_0^{LF}\left[u(w_0^{LF}) - u(b)\right]. \tag{21}$$

Proof. See Appendix A.3 □

The Equal Pay equilibrium is fully recalculated after the policy is imposed: wages, vacancies, market tightness, and job-finding probabilities all adjust endogenously. The relevant maintained assumption is not that vacancy creation fails to adjust by construction, but that Equal Pay does not mechanically eliminate the reduced-form segmentation of search opportunities faced by protected workers. The policy constrains the wage margin, while protected workers remain allocated across the search opportunities available to them in the reduced-form environment.

Equal Pay restricts wage differences directly, but it may also remove the disutility protected-worker submarket from the set of active search options. Protected workers are then redirected toward the non-disutility protected-worker submarket. Non-disutility firms may expand vacancy creation in response, but because vacancy creation is convex, this expansion need not fully replace the vacancies lost when the disutility submarket becomes inactive. Equal Pay therefore generates two opposing forces: a wage-pooling effect and a congestion effect. The wage-pooling effect can raise the wage received by protected workers relative to some *laissez-faire* allocations, while the congestion effect can reduce their job-finding probability. Equal Pay reduces protected-worker welfare precisely when the probability loss dominates the wage gain.

This mechanism should be distinguished from a different policy experiment in which Equal Pay also fully integrates protected and unprotected workers into a single common search market. In such a model, the disappearance of the disutility submarket could raise protected workers' effective outside option to the unprotected-worker level, and non-disutility firms might then treat both groups symmetrically. That is not the experiment studied here. The present analysis isolates the effect of wage regulation when the wage margin is constrained but search opportunities remain segmented in reduced form.[2]

The disappearance of the disutility protected-worker submarket should not be interpreted as affecting only protected workers. In the Equal Pay equilibrium, the removal

[2] A model in which all firms hire both worker types and equal-pay rules are imposed within firms would be needed to study the full market-integration channel. The present analysis should therefore be interpreted as a tractable aggregate benchmark of wage regulation under residual search segmentation.

of disutility vacancies reduces the set of active vacancies available in the economy. Because the expansion of non-disutility vacancies is not infinitely elastic, the policy can also affect unprotected workers through equilibrium vacancy creation, market tightness, and the common wage. The wage-pooling effect may reduce the wage component received by unprotected workers, while the congestion effect operates through the equilibrium job-finding probabilities p_N^{EP} and p_0^{EP}. Thus, the policy cost is not imposed only on protected workers; it can also be borne by unprotected workers through lower job-finding probabilities or a lower wage component.

3.4 Quota Policy Equilibrium

The Quota Policy is modeled as a stylized aggregate representation of firm-level hiring or representation constraints. In practice, quotas or affirmative-action constraints are imposed on firms or organizations. In the model, I capture their aggregate implication by restricting differences in job-finding probabilities across groups in the constrained equilibrium. The policy requires $p_N^Q = p_0^Q = p_d^Q$. The quota does not directly impose wage equality. Wages remain determined by the submarket-specific surplus-sharing rule:

$$w_s^Q = b + \beta(\theta_s^Q) S_s, \qquad s \in \mathcal{S}. \tag{22}$$

The strict quota has an important implication for the disutility protected-worker submarket. Since $p_0^Q = p_d^Q$, and since $p(\theta)$ is strictly increasing, the quota implies $\theta_0^Q = \theta_d^Q$. Therefore, $\beta(\theta_0^Q) = \beta(\theta_d^Q)$. But $S_0 > S_d$, so (22) gives $w_0^Q > w_d^Q$. Protected workers would then strictly prefer submarket 0 to submarket d. Hence, in a decentralized equilibrium, the disutility protected-worker submarket receives no applications:

$$A_d^Q = 0, \qquad V_d^Q = 0. \tag{23}$$

The effective quota constraint is therefore imposed between the unprotected-worker submarket and the active protected-worker non-disutility submarket:

$$p_N^Q = p_0^Q. \tag{24}$$

Since $p(\theta)$ is strictly increasing, (24) implies $\theta_N^Q = \theta_0^Q$. Because $S_N = S_0$, the endogenous surplus-sharing rule then implies

$$w_N^Q = w_0^Q.$$

Thus, unlike Equal Pay, the quota does not impose wage equality directly. It equalizes job-finding probabilities directly, and wage equality follows indirectly from the equality of market tightness and the equality of surplus on the two active non-disutility submarkets.

Definition 3. *A Quota equilibrium is a collection*

$$\left\{A_s^Q, V_s^Q, p_s^Q, q_s^Q, w_s^Q\right\}_{s \in \mathcal{S}}$$

such that:

1. *Wages satisfy* (22).
2. *The disutility protected-worker submarket is inactive, as in* (23).
3. *Matching probabilities satisfy* (6) *and* (7) *on active submarkets.*
4. *Applications clear, with* $A_N^Q = \mu$, $A_0^Q = 1 - \mu$, *and* $A_d^Q = 0$.
5. *The effective quota constraint* (24) *holds.*
6. *Vacancy creation satisfies the constrained vacancy-posting conditions associated with* (5) *and the additional constraint* (24).

Protected-worker expected utility under the quota is

$$U_P^Q = u(b) + p_0^Q \left[u(w_0^Q) - u(b) \right] - \delta. \tag{25}$$

Unprotected-worker expected utility under the quota is

$$U_U^Q = u(b) + p_N^Q \left[u(w_N^Q) - u(b) \right] - \delta. \tag{26}$$

Proposition 4. *The Quota Policy reduces protected-worker welfare relative to laissez-faire if and only if*

$$p_0^Q \left[u(w_0^Q) - u(b) \right] < p_0^{LF} \left[u(w_0^{LF}) - u(b) \right]. \tag{27}$$

Proof. See Appendix A.4 □

The comparison between Equal Pay and the Quota Policy highlights the central mechanism of the model. Equal Pay constrains the wage margin directly. Once wages are equalized, disutility firms cannot attract protected workers by offering a lower-wage, higher-probability contract, and the disutility protected-worker submarket becomes inactive. However, Equal Pay does not impose equality of market tightness between the unprotected-worker and protected-worker submarkets. As a result, the common wage can coexist with different job-finding probabilities. In this sense, Equal Pay restricts wage discrimination but may leave hiring discrimination, understood as unequal job-finding probabilities, in place.

The Quota Policy operates through the opposite channel. It constrains job-finding probabilities directly. Since job-finding probabilities are increasing in market tightness, the quota equalizes tightness across the active unprotected-worker and protected-worker non-disutility submarkets. Because these two active submarkets generate the same match surplus, equal tightness implies the same surplus-sharing term and therefore the same wage. Thus, the quota eliminates hiring-probability differences directly and wage differences indirectly. The two policies therefore eliminate different margins by construction, but neither policy is guaranteed to improve welfare. Welfare depends on whether the new equilibrium wage–probability combination dominates the *laissez-faire* combination available to workers.

The quota removes the disutility protected-worker submarket for a different reason than Equal Pay. Equal Pay restricts wages; once wages are equalized, disutility firms may be unable to attract protected workers by offering a lower-wage, higher-probability allocation. The quota instead restricts job-finding probabilities. Since the disutility submarket generates a lower surplus, equal probabilities imply a lower wage in that submarket, making it unattractive to protected workers. Hence both policies can eliminate the disutility protected-worker submarket in the reduced-form submarket-choice equilibrium, but they do so through opposite constraints. Equal Pay constrains wages and lets probabilities adjust. The quota constrains probabilities and, with tightness-dependent surplus sharing, generates wage equality indirectly across active non-disutility submarkets.

4 Closed-form Characterization

This section provides a closed-form characterization under additional functional assumptions. The previous section defines equilibria for a general utility function, a general constant-returns matching function, a general strictly convex vacancy-cost function, and a tightness-dependent surplus-sharing rule. I now impose simple functional forms in order to obtain transparent expressions for equilibrium utilities and policy comparisons. These assumptions are not needed for the general model; they are introduced only to illustrate the mechanism analytically.

First, I assume risk-neutral workers, $u(w) = w$. This assumption makes expected utility linear in wages and job-finding probabilities. Under risk neutrality, $u(w_s) - u(b) = w_s - b$.

Second, I use the matching function

$$\mathcal{M}(A, V) = \frac{AV}{A + V}. \tag{28}$$

This function is increasing, concave, homogeneous of degree one, and satisfies $\mathcal{M}(A, V) \leq \min\{A, V\}$. It therefore delivers job-finding and vacancy-filling probabilities that are bounded between zero and one. Bounded constant-returns matching functions of this type are commonly used in search-and-matching models (e.g., Den Haan, Ramey, and Watson (2000) and the survey by Petrongolo and Pissarides (2001)). Under (28), the job-finding and vacancy-filling probabilities are $p_s = V_s/(A_s + V_s)$ and $q_s = A_s/(A_s + V_s) = 1 - p_s$.

Third, I assume that the worker's share of surplus is equal to the job-finding probability:

$$\beta(\theta_s) = p(\theta_s) = p_s. \tag{29}$$

This assumption is not meant to be structural. It is a simple reduced-form way to capture the idea that workers capture a larger share of the surplus when the market is tighter.

Since $p'(\theta) > 0$, (29) implies $\beta'(\theta) > 0$. It also keeps the algebra tractable, because the same object governs both employment probabilities and the workers' share of the surplus. Under this assumption, the wage in submarket s is

$$w_s = b + p_s S_s.$$

Hence the expected employment rent in submarket s is $p_s(w_s - b) = p_s^2 S_s$.

Finally, I assume quadratic vacancy costs:

$$C(V_s) = \psi V_s + \frac{\kappa}{2} V_s^2, \qquad \psi > 0, \quad \kappa > 0. \tag{30}$$

This is the closed-form counterpart of the general strictly convex vacancy cost introduced in Section 3.1. Convex vacancy or hiring costs are widely used to capture adjustment costs, limited managerial capacity, or decreasing returns in job creation (e.g., Merz (1995), Andolfatto (1996), and Gertler, Sala, and Trigari (2008)). Here, the quadratic term is useful because it prevents vacancy creation from being infinitely elastic while preserving analytical tractability.

4.1 Main closed-form results

This subsection presents the main closed-form implications. The algebra behind these expressions is collected in Appendix B.

Laissez-faire. Suppose both protected-worker submarkets are active under *laissez-faire*. Since $w_s^{LF} - b = p_s^{LF} S_s$, protected-worker expected utility in submarket $s \in \{0, d\}$ is $b + (p_s^{LF})^2 S_s - \delta$. Submarket-choice indifference therefore implies

$$(p_0^{LF})^2 S_0 = (p_d^{LF})^2 S_d. \tag{31}$$

Thus,

$$p_d^{LF} = p_0^{LF} \sqrt{\frac{S_0}{S_d}}. \tag{32}$$

Since $S_0 > S_d$, one has $p_d^{LF} > p_0^{LF}$. The disutility submarket compensates protected workers for its lower surplus by offering a higher job-finding probability. Protected-worker expected utility under *laissez-faire* is

$$U_P^{LF} = b + (p_0^{LF})^2 S_0 - \delta = b + (p_d^{LF})^2 S_d - \delta. \tag{33}$$

Under the functional assumptions above, vacancy creation in protected-worker submarket $s \in \{0, d\}$ satisfies

$$V_s^{LF} = \frac{(1 - p_s^{LF})^2 S_s - \psi}{\kappa}. \tag{34}$$

Together with (32) and application clearing, equation (34) pins down p_0^{LF}, and therefore the entire protected-worker equilibrium.

The disutility submarket is active if and only if

$$0 < p_0^{LF} < \sqrt{\frac{S_d}{S_0}} \tag{35}$$

and

$$V_d^{LF} > 0. \tag{36}$$

The first condition ensures that $p_d^{LF} \in (0,1)$. The second condition ensures that disutility firms create vacancies. Under the maintained interior-participation assumption, protected workers also require $(p_0^{LF})^2 S_0 \geq \delta$.

Equal Pay. Equal Pay may eliminate the disutility submarket when disutility firms cannot profitably compensate their cost by offering higher job-finding probabilities, hence

$$A_d^{EP} = V_d^{EP} = 0.$$

The reason is that Equal Pay removes the wage difference between the two protected-worker submarkets. If the disutility submarket offered a lower job-finding probability, no protected worker would apply to it; if it tried to offer the same probability, disutility firms would be less profitable because they still bear the cost d. Hence the relevant Equal Pay equilibrium has a single active protected-worker submarket, namely the non-disutility one. Therefore $A_0^{EP} = 1 - \mu$.

The Equal Pay wage is not fixed at the *laissez-faire* unprotected-worker wage. It is recalculated in equilibrium. Since only the unprotected-worker submarket and the non-disutility protected-worker submarket are active, and since $S_N = S_0$, the common Equal Pay wage satisfies

$$w^{EP} - b = S_0 \frac{p_N^{EP} \mathcal{M}(A_N^{EP}, V_N^{EP}) + p_0^{EP} \mathcal{M}(A_0^{EP}, V_0^{EP})}{\mathcal{M}(A_N^{EP}, V_N^{EP}) + \mathcal{M}(A_0^{EP}, V_0^{EP})}. \tag{37}$$

This expression follows from $\beta(\theta_s) = p_s$. The common wage is therefore a match-weighted average of the surplus shares generated in the two active non-disutility submarkets.

Protected-worker expected utility under Equal Pay is

$$U_P^{EP} = b + p_0^{EP}(w^{EP} - b) - \delta. \tag{38}$$

Equal Pay reduces protected-worker welfare if and only if

$$p_0^{EP}(w^{EP} - b) < (p_0^{LF})^2 S_0.$$

Equivalently, Equal Pay is costly for protected workers whenever

$$p_0^{EP} < \frac{(p_0^{LF})^2 S_0}{w^{EP} - b}. \tag{39}$$

This is the explicit probability threshold. Equal Pay may raise the wage received by protected workers relative to some *laissez-faire* allocations, but it reduces protected-worker welfare whenever the equilibrium job-finding probability in the remaining protected-worker submarket falls sufficiently.

The disappearance of the disutility submarket also affects the aggregate allocation of vacancies. Non-disutility firms may create additional vacancies after Equal Pay is imposed, but because vacancy creation is convex, this response is not infinitely elastic and need not fully offset the lost disutility vacancies. The Equal Pay effect should therefore be interpreted as the combination of a wage-pooling effect and a congestion effect. The latter operates through the post-policy job-finding probabilities of both protected and unprotected workers.

Quota Policy. Consider now a strict Quota Policy. The policy imposes equality of job-finding probabilities across unprotected and protected workers in the new constrained equilibrium. It does not impose the *laissez-faire* unprotected-worker probability. The relevant probability is recalculated in equilibrium.

In a decentralized submarket-choice equilibrium, the disutility protected-worker submarket is again inactive:

$$A_d^Q = V_d^Q = 0.$$

The reason is different from Equal Pay. Under a strict quota, the policy equalizes job-finding probabilities. Since $p_0^Q = p_d^Q$, and since $p(\theta)$ is strictly increasing, one has $\theta_0^Q = \theta_d^Q$. Hence $\beta(\theta_0^Q) = \beta(\theta_d^Q)$. But $S_0 > S_d$, so $w_0^Q > w_d^Q$. Protected workers therefore strictly prefer the non-disutility submarket, and the disutility submarket receives no applications.

The effective quota constraint is therefore imposed between the unprotected-worker submarket and the active non-disutility protected-worker submarket:

$$p_N^Q = p_0^Q. \tag{40}$$

Since $p(\theta)$ is strictly increasing, this implies $\theta_N^Q = \theta_0^Q$. Because $S_N = S_0$, and because $\beta(\theta_N^Q) = \beta(\theta_0^Q)$, the quota also implies $w_N^Q = w_0^Q$ indirectly. Thus, the quota equalizes probabilities directly and wages indirectly. In the closed-form case, let $p^Q = p_N^Q = p_0^Q$. Then

$$w_0^Q - b = p^Q S_0.$$

Protected-worker expected utility under the quota is

$$U_P^Q = b + (p^Q)^2 S_0 - \delta. \tag{41}$$

The quota reduces protected-worker welfare if and only if

$$(p^Q)^2 S_0 < (p_0^{LF})^2 S_0.$$

Since $S_0 > 0$, this is equivalent to

$$p^Q < p_0^{LF}. \tag{42}$$

Thus, the strict quota is costly for protected workers whenever the common job-finding probability induced by the policy is lower than the job-finding probability that protected workers obtained in the non-disutility submarket under *laissez-faire*. With tightness-dependent surplus sharing, this probability decline also lowers the wage, because $w_0^Q - b = p^Q S_0$.

5 Aggregate Welfare and Policy Incidence

The previous sections characterize the individual welfare effects of Equal Pay and Quota Policies. This section asks a different question: who bears the welfare cost of each policy, and when does the policy reduce aggregate welfare? To obtain sharp conditions, I use the closed-form assumptions introduced in Section 4. In particular, workers are risk neutral, $u(w) = w$, the worker's surplus share satisfies $\beta(\theta_s) = p_s$, and the active non-disutility submarkets generate the same surplus $S_N = S_0 = y - b$.

The total mass of workers is normalized to one. The mass of unprotected workers is $\mu \in (0,1)$, and the mass of protected workers is $1 - \mu$. For any policy regime $\pi \in \{LF, EP, Q\}$, aggregate welfare is evaluated using the ex-ante utilitarian criterion

$$W^\pi = \mu U_U^\pi + (1-\mu) U_P^\pi.$$

5.1 Aggregate welfare

Under *laissez-faire*, expected utilities are

$$U_U^{LF} = b + (p_N^{LF})^2 S_0 - \delta,$$

and

$$U_P^{LF} = b + (p_0^{LF})^2 S_0 - \delta.$$

The second expression uses the protected-worker submarket associated with non-disutility firms. Since protected workers are indifferent across active protected-worker submarkets under *laissez-faire*, the same expected utility is obtained using the disutility submarket. Therefore,

$$W^{LF} = b - \delta + S_0 \left[\mu (p_N^{LF})^2 + (1-\mu)(p_0^{LF})^2\right].$$

Proposition 5. *Under the closed-form assumptions, Equal Pay reduces aggregate welfare if and only if*

$$\mu (p_N^{EP})^2 + (1-\mu)(p_0^{EP})^2 < \mu (p_N^{LF})^2 + (1-\mu)(p_0^{LF})^2.$$

Proof. See Appendix A.5 □

Proposition 5 shows that Equal Pay reduces aggregate welfare whenever the policy-induced equilibrium lowers the population-weighted sum of squared job-finding probabilities. The relevant comparison is not only whether the common wage rises or falls, but whether the wage gain generated by Equal Pay is sufficient to compensate for the induced change in group-specific job-finding probabilities. Since the common Equal Pay wage is itself determined by equilibrium matching outcomes, aggregate welfare can be written entirely in terms of the job-finding probabilities generated by the policy equilibrium. This condition also makes clear that the cost of Equal Pay need not be borne only by protected workers: both the unprotected-worker probability p_N^{EP} and the protected-worker probability p_0^{EP} enter aggregate welfare.

I now turn to the Quota Policy. Under the strict quota, the active unprotected-worker and protected-worker non-disutility submarkets share a common job-finding probability:

$$p_N^Q = p_0^Q \equiv p^Q.$$

This probability is not inherited from the *laissez-faire* equilibrium. It is the common probability that emerges after firms and workers re-optimize under the quota constraint. Since the quota also implies a common wage in the closed-form environment, expected utilities satisfy

$$U_U^Q = U_P^Q = b + (p^Q)^2 S_0 - \delta.$$

Therefore,

$$W^Q = b - \delta + (p^Q)^2 S_0.$$

For the aggregate welfare comparison under the quota, define the *laissez-faire* root-mean-square job-finding probability as

$$p_{\text{RMS}}^{LF} \equiv \sqrt{\mu(p_N^{LF})^2 + (1-\mu)(p_0^{LF})^2}.$$

This object is the relevant *laissez-faire* benchmark because, under the closed-form assumptions, expected employment rents are quadratic in job-finding probabilities.

Proposition 6. *Under the closed-form assumptions, the Quota Policy reduces aggregate welfare if and only if*

$$p^Q < p_{\text{RMS}}^{LF}.$$

Proof. See Appendix A.6 □

The relevant aggregate benchmark is therefore not the average *laissez-faire* job-finding probability, but the root-mean-square of *laissez-faire* probabilities. This is because, under the closed-form assumptions, expected employment rents are quadratic in job-finding probabilities. The quota is welfare-reducing whenever the common probability generated by the constrained equilibrium falls below this quadratic benchmark.

5.2 Policy incidence

The aggregate welfare conditions above do not yet identify who bears the cost of each policy. This subsection characterizes policy incidence across unprotected and protected workers. Equal Pay is a wage-pooling policy, but not an allocation-pooling policy. It imposes a common wage while allowing job-finding probabilities to remain group-specific. Define

$$\bar{p}^{EP} \equiv \frac{\mu(p_N^{EP})^2 + (1-\mu)(p_0^{EP})^2}{\mu p_N^{EP} + (1-\mu)p_0^{EP}}.$$

Then the Equal Pay wage satisfies

$$w^{EP} - b = S_0 \bar{p}^{EP}.$$

Therefore,

$$U_U^{EP} = b + S_0 p_N^{EP} \bar{p}^{EP} - \delta,$$

and

$$U_P^{EP} = b + S_0 p_0^{EP} \bar{p}^{EP} - \delta.$$

Proposition 7. *Under Equal Pay, group $g \in \{N, 0\}$ loses relative to laissez-faire if and only if*

$$p_g^{EP} \bar{p}^{EP} < (p_g^{LF})^2,$$

where group N denotes unprotected workers and group 0 denotes protected workers in the non-disutility submarket.

Proof. See Appendix A.7 □

Proposition 7 shows that the incidence of Equal Pay depends jointly on two objects: the common wage component, summarized by $\bar{p}^{EP}$, and the group-specific job-finding probability generated by the Equal Pay equilibrium. A group bears the cost of Equal Pay whenever the product of its policy job-finding probability and the common wage component falls below its *laissez-faire* employment rent. Thus, although Equal Pay directly restricts the wage margin, its incidence depends on how the policy reshapes both wage pooling and congestion across groups.

The Quota Policy has a different incidence structure. It is an allocation-pooling policy. It imposes a common job-finding probability and, in the closed-form environment, generates a common wage. Each group receives

$$U^Q = b + (p^Q)^2 S_0 - \delta.$$

Hence,

$$\Delta U_U^Q < 0 \iff p^Q < p_N^{LF},$$

and

$$\Delta U_P^Q < 0 \iff p^Q < p_0^{LF}.$$

The following proposition gives a complete incidence classification for the natural case in which unprotected workers have a higher *laissez-faire* job-finding probability than protected workers in the non-disutility submarket.

Proposition 8. *Suppose $p_N^{LF} > p_0^{LF}$. Under the Quota Policy, incidence is characterized by the position of p^Q relative to p_0^{LF}, p_{RMS}^{LF}, and p_N^{LF}. In particular,*

- *if $p^Q < p_0^{LF}$, both groups lose and aggregate welfare falls;*
- *if $p_0^{LF} \leq p^Q < p_{\mathrm{RMS}}^{LF}$, protected workers gain, unprotected workers lose, and aggregate welfare falls;*
- *if $p_{\mathrm{RMS}}^{LF} \leq p^Q < p_N^{LF}$, protected workers gain, unprotected workers lose, and aggregate welfare rises; and*
- *if $p^Q \geq p_N^{LF}$, both groups gain and aggregate welfare rises.*

Proof. See Appendix A.8 □

The Quota Policy therefore replaces the two *laissez-faire* employment rents, $(p_N^{LF})^2 S_0$ and $(p_0^{LF})^2 S_0$, with a common rent, $(p^Q)^2 S_0$. When the common probability lies below p_0^{LF}, both groups are worse off. When it lies between p_0^{LF} and p_N^{LF}, the quota redistributes expected utility from unprotected workers to protected workers. Whether this redistribution increases or decreases aggregate welfare depends on whether p^Q lies above or below the quadratic benchmark p_{RMS}^{LF}. Finally, if the quota generates a common probability above p_N^{LF}, both groups gain.

The key distinction is therefore that Equal Pay pools wages, whereas the Quota Policy pools allocations. Equal Pay imposes a common wage but leaves job-finding probabilities group-specific. The quota imposes a common probability and, in this environment, also generates a common wage. As a result, Equal Pay incidence depends on the interaction between the common wage component and group-specific congestion, whereas quota incidence depends only on the position of the common probability relative to each group's *laissez-faire* probability.

6 Discussion and Conclusion

This paper studies how anti-discrimination policies affect welfare when firms and workers can adjust across wage and hiring-probability margins. The analysis is conducted in a reduced-form search environment with submarket choice, endogenous vacancy creation, taste-based disutility, and surplus-sharing wage determination. The model is not a full wage-posting, bargaining, or competitive-search equilibrium. Instead, it provides a tractable benchmark for studying how policy constraints reshape the wage–employment-probability allocation.

The main lesson is that the welfare effect of an anti-discrimination policy cannot be inferred from its success on the targeted margin alone. Equal Pay restricts wage differences, but it may shift adjustment toward employment probabilities when vacancy creation is not sufficiently elastic. In the model, the disutility protected-worker submarket may become inactive. Protected workers then apply through the non-disutility protected-worker submarket, while non-disutility firms may create additional vacancies. Because vacancy creation is convex, this expansion is not infinitely elastic and need not fully offset the loss of disutility vacancies. Equal Pay therefore generates both a wage effect and a congestion effect. The wage effect follows from pooling wages across active submarkets, while the congestion effect operates through the equilibrium job-finding probabilities of protected and unprotected workers. As a result, Equal Pay may raise wages conditional on employment but reduce expected welfare for one or both groups.

The Quota Policy operates through the opposite margin. It restricts differences in job-finding probabilities and, in the closed-form version of the model, indirectly equalizes wages across active non-disutility submarkets. However, this equality may be achieved at a common allocation that is less favorable than the *laissez-faire* allocation for one or both groups. Equal Pay and Quota Policies therefore have different incidence structures. Equal Pay pools wages while leaving probabilities group-specific. The quota pools the wage–probability allocation.

These results do not imply that anti-discrimination policy is undesirable. Rather, they show that policy evaluation must account for equilibrium adjustment. A policy that eliminates an observed wage gap may fail to improve welfare if it worsens employment access. A policy that equalizes hiring probabilities may still reduce welfare if the common allocation is sufficiently low. The relevant object for welfare analysis is therefore the full expected employment rent, not only the margin directly targeted by regulation.

The analysis also has important limitations. First, the surplus-sharing rule is assumed rather than derived from a full wage-posting or bargaining game. It should therefore be read as a reduced-form wage rule, not as the equilibrium outcome of a fully specified wage-setting process. Second, the policy constraints are modeled as aggregate reduced-form restrictions, whereas actual equal-pay and quota policies are imposed at the firm or organization level. The Equal Pay and Quota exercises should therefore be interpreted as stylized aggregate representations of firm-level legal constraints, not as literal descriptions of legal implementation within firms. Third, the segmentation between protected and unprotected workers is imposed as a reduced-form barrier rather than derived endogenously. This segmentation captures institutional, informational, occupational, spatial, or network frictions that prevent full arbitrage across all vacancies. These assumptions make the mechanism transparent, but they also delimit the interpretation of the results.

Future research could relax these assumptions. One extension would solve a full wage-posting, bargaining, or competitive-search equilibrium with taste-based disutility. Another would model firm-level equal-pay and quota constraints explicitly, allowing all firms to hire both worker types and to offer type-specific wages subject to legal restrictions. A further extension would allow discrimination to move into job quality, hours, promotion opportunities, occupational assignment, or other non-wage margins. Such extensions would show whether the substitution mechanism identified here survives in richer environments.

The paper's contribution is therefore cautionary but constructive. Anti-discrimination policies should not be evaluated only by whether they eliminate a targeted inequality. They should be evaluated by how they reshape wages, vacancies, job-finding probabilities, and welfare jointly. When discriminatory practices are substitutable, effective policy design requires accounting for the full set of behavioral responses generated by regulation.

References

Andolfatto, David (1996). "Business cycles and labor-market search". In: *The american economic review*, pp. 112–132.

Arrow, Kenneth (1971). "The theory of discrimination". In: *OUP Oxford.*

— (1972). "Some mathematical models of race discrimination in the labor market". In: *Racial discrimination in economic life*, pp. 187–204.

Arrow, Kenneth and Gerard Debreu (1954). "Existence of an equilibrium for a competitive economy". In: *The Foundations of Price Theory Vol 5*. Routledge, pp. 289–316.

Bartos, Vojtech et al. (2013). "Attention discrimination: Theory and field experiments". In: *CERGE-EI Working Paper Series* 499.

Becker, Gary S (1957). *The economics of discrimination*. University of Chicago press.

Bertrand, Marianne and Esther Duflo (2017). "Field experiments on discrimination". In: *Handbook of economic field experiments* 1, pp. 309–393.

Bertrand, Marianne and Sendhil Mullainathan (2004). "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination". In: *American economic review* 94.4, pp. 991–1013.

Black, Dan A (1995). "Discrimination in an equilibrium search model". In: *Journal of labor Economics* 13.2, pp. 309–334.

Blau, Francine D and Lawrence M Kahn (2017). "The gender wage gap: Extent, trends, and explanations". In: *Journal of economic literature* 55.3, pp. 789–865.

Bohren, J Aislinn, Alex Imas, and Michael Rosenberg (2019). "The dynamics of discrimination: Theory and evidence". In: *American economic review* 109.10, pp. 3395–3436.

Bordalo, Pedro (2016). "Beliefs about Gender Pedro Bordalo, Katie Coffman, Nicola Gennaioli, and Andrei Shleifer December 2016". In: *Harvard Business School.*

Bowlus, Audra J and Zvi Eckstein (2002). "Discrimination and skill differences in an equilibrium search model". In: *International Economic Review* 43.4, pp. 1309–1345.

Burdett, Kenneth and Dale T Mortensen (1998). "Wage differentials, employer size, and unemployment". In: *International economic review*, pp. 257–273.

Calvo-Armengol, Antoni and Matthew O Jackson (2004). "The effects of social networks on employment and inequality". In: *American economic review* 94.3, pp. 426–454.

Charles, Kerwin Kofi and Jonathan Guryan (2011). "Studying discrimination: Fundamental challenges and recent progress". In: *Annu. Rev. Econ.* 3.1, pp. 479–511.

Council, National Research et al. (1981). "Women, work, and wages: Equal pay for jobs of equal value". In: *The National Academies Press.*

Den Haan, Wouter J, Garey Ramey, and Joel Watson (2000). "Job destruction and propagation of shocks". In: *American economic review* 90.3, pp. 482–498.

Diamond, Peter A (1971). "A model of price adjustment". In: *Journal of economic theory* 3.2, pp. 156–168.

— (1982). "Aggregate demand management in search equilibrium". In: *Journal of political Economy* 90.5, pp. 881–894.

Donohue, John J (2005). *The law and economics of antidiscrimination law.*

Donohue, John J and James J Heckman (1991). *Continuous versus episodic change: The impact of civil rights policy on the economic status of blacks.*

European Union, Council of the (2000a). *Council Directive 2000/43/EC of 29 June 2000 implementing the principle of equal treatment between persons irrespective of racial or ethnic origin*. Accessed: 2025-02-13. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32000L0043.

— (2000b). *Directive 2000/78/EC establishing a general framework for equal treatment in employment and occupation*. Accessed: 2024-11-20. URL: https://eur-lex.europa.eu/legal-content/FR/TXT/?uri=CELEX:32000L0078.

Flabbi, Luca (2010). "Gender discrimination estimation in a search model with matching and bargaining". In: *International Economic Review* 51.3, pp. 745–783.

Galos, Diana Roxana and Alexander Coppock (2023). "Gender composition predicts gender bias: A meta-reanalysis of hiring discrimination audit experiments". In: *Science Advances* 9.18, eade7979.

Gelbach, Jonah B (2016). "When do covariates matter? And which ones, and how much?" In: *Journal of Labor Economics* 34.2, pp. 509–543.

Gertler, Mark, Luca Sala, and Antonella Trigari (2008). "An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining". In: *Journal of Money, Credit and Banking* 40.8, pp. 1713–1764.

Heckman, James J and Brook S Payner (1989). *Determining the impact of federal antidiscrimination policy on the economic status of blacks: A study of South Carolina*.

Heckman, James J and Peter Siegelman (1993). "The Urban Institute Audit Studies: Their Methods and Findings". In: *Clear and Convincing Evidence: Measurement of Discrimination in America*.

Holzer, Harry and David Neumark (2000). "Assessing affirmative action". In: *Journal of Economic literature* 38.3, pp. 483–568.

Kaas, Leo and Jun Lu (2010). "Equal-treatment policy in a random search model with taste discrimination". In: *Labour Economics* 17.4, pp. 699–709.

Lang, Kevin and Jee-Yeon K Lehmann (2012). "Racial discrimination in the labor market: Theory and empirics". In: *Journal of Economic Literature* 50.4, pp. 959–1006.

Lang, Kevin, Michael Manove, and William T Dickens (2005). "Racial discrimination in labor markets with posted wage offers". In: *American Economic Review* 95.4, pp. 1327–1340.

Merz, Monika (1995). "Search in the labor market and the real business cycle". In: *Journal of monetary Economics* 36.2, pp. 269–300.

Moen, Espen R (1997). "Competitive search equilibrium". In: *Journal of political Economy* 105.2, pp. 385–411.

Mortensen, Dale T and Christopher A Pissarides (1994). "Job creation and job destruction in the theory of unemployment". In: *The review of economic studies* 61.3, pp. 397–415.

Neumark, David (2018). "Experimental research on labor market discrimination". In: *Journal of Economic Literature* 56.3, pp. 799–866.

Oaxaca, Ronald (1973). "Male-female wage differentials in urban labor markets". In: *International economic review*, pp. 693–709.

Pager, Devah (2008). *Marked: Race, crime, and finding work in an era of mass incarceration*. University of Chicago Press.

Pager, Devah and Hana Shepherd (2008). "The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets". In: *Annu. Rev. Sociol* 34.1, pp. 181–209.

Petrongolo, Barbara and Christopher A Pissarides (2001). "Looking into the black box: A survey of the matching function". In: *Journal of Economic literature* 39.2, pp. 390–431.

Phelps, Edmund S (1972). "The statistical theory of racism and sexism". In: *The american economic review* 62.4, pp. 659–661.

Pissarides, Christopher A (2000). *Equilibrium unemployment theory*. MIT press.

Polachek, Solomon W (2019). "Equal pay legislation and the gender wage gap". In: *IZA World of Labor*.

Reskin, Barbara F (2000). "The proximate causes of employment discrimination". In: *Contemporary sociology* 29.2, pp. 319–328.

Rogerson, Richard, Robert Shimer, and Randall Wright (2005). "Search-theoretic models of the labor market: A survey". In: *Journal of economic literature* 43.4, pp. 959–988.

Rosen, Sherwin (1974). "Hedonic prices and implicit markets: product differentiation in pure competition". In: *Journal of political economy* 82.1, pp. 34–55.

Walras, Léon (1877). *Éléments d'économie politique pure: ou, Théorie de la richesse sociale*. F. Rouge.

Ziller, Conrad and Marc Helbling (2019). "Antidiscrimination laws, policy knowledge and political support". In: *British Journal of Political Science* 49.3, pp. 1027–1044.

A Proofs

A.1 Proof Proposition 1

By (9), the condition $\beta(\theta_0^{LF})S_0 > \beta(\theta_d^{LF})S_d$ implies $w_0^{LF} > w_d^{LF}$. Since u is strictly increasing, $u(w_0^{LF}) - u(b) > u(w_d^{LF}) - u(b)$. If both protected-worker submarkets are active, protected workers must be indifferent between them, so (10) must hold. The equality can hold only if $p_0^{LF} < p_d^{LF}$.

A.2 Proof Proposition 2

Under Equal Pay, protected workers receive the same wage in the disutility and non-disutility protected-worker submarkets: $w_d^{EP} = w_0^{EP} = w^{EP}$. Hence protected workers rank the two protected-worker submarkets only by their job-finding probabilities. Since $p(\theta)$ is strictly increasing, this is equivalent to ranking submarkets by tightness.

Suppose first that $\theta_d^{EP} < \theta_0^{EP}$. Then $p_d^{EP} = p(\theta_d^{EP}) < p(\theta_0^{EP}) = p_0^{EP}$. Since wages are equal, protected workers strictly prefer submarket 0 to submarket d. Therefore $A_d^{EP} = 0$.

Suppose next that $\theta_d^{EP} = \theta_0^{EP}$. Then $p_d^{EP} = p_0^{EP}$. Protected workers are indifferent between the two protected-worker submarkets. By the tie-breaking convention, they apply to the non-disutility protected-worker submarket. Hence again $A_d^{EP} = 0$.

It remains to consider whether disutility firms could attract protected workers by offering a strictly higher job-finding probability. This would require $p_d^{EP} > p_0^{EP}$. Since $p(\theta)$ is strictly increasing, this implies $\theta_d^{EP} > \theta_0^{EP}$. Because the vacancy-filling probability $q(\theta)$ is strictly decreasing in tightness, it follows that $q(\theta_d^{EP}) < q(\theta_0^{EP})$. For a disutility firm, the net surplus from a filled vacancy under Equal Pay is $y - d - w^{EP}$. Thus, for any disutility submarket tightness satisfying $\theta_d^{EP} \geq \theta_0^{EP}$, the expected return from a marginal disutility vacancy is bounded above by

$$q(\theta_d^{EP})(y - d - w^{EP}) \leq q(\theta_0^{EP})(y - d - w^{EP}).$$

By condition (16),

$$q(\theta_0^{EP})(y - d - w^{EP}) \leq C'(0).$$

Therefore,

$$q(\theta_d^{EP})(y - d - w^{EP}) \leq C'(0).$$

Since the vacancy cost is strictly convex, $C'(V) \geq C'(0)$ for all $V \geq 0$. Hence, for any $V_d^{EP} > 0$,

$$q(\theta_d^{EP})(y - d - w^{EP}) < C'(V_d^{EP})$$

whenever C' is strictly increasing, and weakly so at the boundary. Thus the vacancy-creation condition

$$C'(V_d^{EP}) = q(\theta_d^{EP})(y - d - w^{EP})$$

cannot hold for any positive mass of disutility vacancies capable of attracting protected workers away from submarket 0. Therefore disutility firms cannot profitably create vacancies in a submarket with $p_d^{EP} > p_0^{EP}$.

Hence no active disutility protected-worker submarket can be sustained in the relevant Equal Pay equilibrium. Therefore $A_d^{EP} = 0$. Given $A_d^{EP} = 0$, the vacancy-filling probability for disutility vacancies is zero. Opening vacancies in submarket d then yields no expected hiring benefit and only entails vacancy-creation costs. The vacancy-creation condition therefore implies $V_d^{EP} = 0$.

A.3 Proof Proposition 3

Suppose first that Equal Pay reduces protected-worker welfare. Then $U_P^{EP} < U_P^{LF}$. Using (19) and (11), this means

$$u(b) + p_0^{EP}\left[u(w^{EP}) - u(b)\right] - \delta < u(b) + p_0^{LF}\left[u(w_0^{LF}) - u(b)\right] - \delta.$$

Canceling $u(b)$ and δ gives (21).

Conversely, suppose (21) holds. Adding $u(b) - \delta$ to both sides gives $U_P^{EP} < U_P^{LF}$. Therefore, the condition is both necessary and sufficient.

A.4 Proof Proposition 4

Suppose first that the quota reduces protected-worker welfare. Then $U_P^{Q} < U_P^{LF}$. Using (25) and (11), this means

$$u(b) + p_0^{Q}\left[u(w_0^{Q}) - u(b)\right] - \delta < u(b) + p_0^{LF}\left[u(w_0^{LF}) - u(b)\right] - \delta.$$

Canceling $u(b)$ and δ gives (27).

Conversely, suppose (27) holds. Adding $u(b) - \delta$ to both sides gives $U_P^{Q} < U_P^{LF}$. Therefore, the condition is both necessary and sufficient.

A.5 Proof Proposition 5

As established in Proposition 2, under the tie-breaking convention and the no-profitable-disutility-entry condition, the disutility protected-worker submarket is inactive in the relevant Equal Pay equilibrium. The active submarkets are the unprotected-worker submarket N and the protected-worker submarket associated with non-disutility firms, denoted by 0. Expected utilities are

$$U_U^{EP} = b + p_N^{EP}(w^{EP} - b) - \delta,$$

and

$$U_P^{EP} = b + p_0^{EP}(w^{EP} - b) - \delta.$$

Hence,

$$W^{EP} = b - \delta + (w^{EP} - b)\left[\mu p_N^{EP} + (1-\mu)p_0^{EP}\right].$$

From the Equal Pay wage rule in the closed-form environment,

$$w^{EP} - b = S_0 \frac{\mu(p_N^{EP})^2 + (1-\mu)(p_0^{EP})^2}{\mu p_N^{EP} + (1-\mu)p_0^{EP}}.$$

Substituting this expression into W^{EP} gives

$$W^{EP} = b - \delta + S_0 \left[\mu(p_N^{EP})^2 + (1-\mu)(p_0^{EP})^2\right].$$

Comparing this expression with W^{LF}, and using $S_0 > 0$, yields the stated condition.

A.6 Proof Proposition 6

The Quota Policy reduces aggregate welfare if and only if

$$W^Q < W^{LF}.$$

Using the expressions above, this is equivalent to

$$b - \delta + (p^Q)^2 S_0 < b - \delta + S_0 \left[\mu(p_N^{LF})^2 + (1-\mu)(p_0^{LF})^2\right].$$

Since $S_0 > 0$, this becomes

$$(p^Q)^2 < \mu(p_N^{LF})^2 + (1-\mu)(p_0^{LF})^2.$$

Taking square roots gives

$$p^Q < p_{\text{RMS}}^{LF}.$$

A.7 Proof Proposition 7

For unprotected workers,

$$\Delta U_U^{EP} = U_U^{EP} - U_U^{LF} = S_0 \left[p_N^{EP}\bar{p}^{EP} - (p_N^{LF})^2\right].$$

Since $S_0 > 0$, unprotected workers lose if and only if

$$p_N^{EP}\bar{p}^{EP} < (p_N^{LF})^2.$$

For protected workers,

$$\Delta U_P^{EP} = U_P^{EP} - U_P^{LF} = S_0 \left[p_0^{EP}\bar{p}^{EP} - (p_0^{LF})^2\right].$$

Since $S_0 > 0$, protected workers lose if and only if

$$p_0^{EP}\bar{p}^{EP} < (p_0^{LF})^2.$$

This gives the stated group-specific condition.

A.8 Proof Proposition 8

Under the quota, both groups receive the common utility

$$U^Q = b + (p^Q)^2 S_0 - \delta.$$

Unprotected workers lose if and only if

$$b + (p^Q)^2 S_0 - \delta < b + (p_N^{LF})^2 S_0 - \delta,$$

which is equivalent to

$$p^Q < p_N^{LF}.$$

Protected workers lose if and only if

$$b + (p^Q)^2 S_0 - \delta < b + (p_0^{LF})^2 S_0 - \delta,$$

which is equivalent to

$$p^Q < p_0^{LF}.$$

Aggregate welfare falls if and only if

$$p^Q < p_{\mathrm{RMS}}^{LF},$$

by Proposition 6. Since $p_N^{LF} > p_0^{LF}$ and both groups have positive mass, the root-mean-square probability satisfies

$$p_0^{LF} < p_{\mathrm{RMS}}^{LF} < p_N^{LF}.$$

Combining these three thresholds yields the four regions stated in the proposition.

B Computational details for the closed-form equilibrium

This section provides the algebra underlying the closed-form expressions used in the closed-form characterization.

Laissez-faire. Under *laissez-faire*, protected-worker indifference gives (31). Therefore,

$$p_d^{LF} = p_0^{LF} \sqrt{\frac{S_0}{S_d}}.$$

Since $q_s^{LF} = 1 - p_s^{LF}$, and since $w_s^{LF} - b = p_s^{LF} S_s$, the surplus retained by firms per filled vacancy is

$$y - c_s - w_s^{LF} = S_s - (w_s^{LF} - b) = (1 - p_s^{LF}) S_s.$$

The vacancy-creation condition is therefore

$$\psi + \kappa V_s^{LF} = (1 - p_s^{LF})^2 S_s, \qquad s \in \{0, d\}.$$

This yields (34).

Application clearing pins down p_0^{LF}. Under (28), $p_s^{LF} = V_s^{LF}/(A_s^{LF} + V_s^{LF})$, so $A_s^{LF} = V_s^{LF}(1 - p_s^{LF})/p_s^{LF}$. Since $A_0^{LF} + A_d^{LF} = 1 - \mu$, one obtains

$$1 - \mu = V_0^{LF} \frac{1 - p_0^{LF}}{p_0^{LF}} + V_d^{LF} \frac{1 - p_0^{LF}\sqrt{\frac{S_0}{S_d}}}{p_0^{LF}\sqrt{\frac{S_0}{S_d}}}. \tag{43}$$

Substituting (34) and (32) into (43) gives a single equation in p_0^{LF}:

$$1 - \mu = \frac{\left[(1 - p_0^{LF})^2 S_0 - \psi\right](1 - p_0^{LF})}{\kappa p_0^{LF}} + \frac{\left[\left(1 - p_0^{LF}\sqrt{\frac{S_0}{S_d}}\right)^2 S_d - \psi\right]\left(1 - p_0^{LF}\sqrt{\frac{S_0}{S_d}}\right)}{\kappa p_0^{LF}\sqrt{\frac{S_0}{S_d}}}. \tag{44}$$

The economically relevant solution is the root satisfying (35), (36), and the interior-participation condition $(p_0^{LF})^2 S_0 \geq \delta$. Once p_0^{LF} is known, p_d^{LF}, V_0^{LF}, V_d^{LF}, and U_P^{LF} follow directly from (32), (34), and (33).

Equal Pay. Under Equal Pay, the disutility protected-worker submarket is inactive, so

$$A_d^{EP} = V_d^{EP} = 0.$$

The active submarkets are the unprotected-worker submarket N and the non-disutility protected-worker submarket 0. With the harmonic matching function, $\mathcal{M}(A_s, V_s) = p_s A_s = q_s V_s$. The Equal Pay wage rule can therefore be written as

$$w^{EP} - b = S_0 \frac{p_N^{EP}\mathcal{M}(A_N^{EP}, V_N^{EP}) + p_0^{EP}\mathcal{M}(A_0^{EP}, V_0^{EP})}{\mathcal{M}(A_N^{EP}, V_N^{EP}) + \mathcal{M}(A_0^{EP}, V_0^{EP})}. \tag{45}$$

Vacancy creation on the active protected-worker submarket satisfies

$$V_0^{EP} = \frac{(1 - p_0^{EP})\left[S_0 - (w^{EP} - b)\right] - \psi}{\kappa}. \tag{46}$$

Vacancy creation on the active unprotected-worker submarket satisfies

$$V_N^{EP} = \frac{(1 - p_N^{EP})\left[S_0 - (w^{EP} - b)\right] - \psi}{\kappa}. \tag{47}$$

Together with $A_N^{EP} = \mu$, $A_0^{EP} = 1 - \mu$, and

$$p_N^{EP} = \frac{V_N^{EP}}{\mu + V_N^{EP}}, \qquad p_0^{EP} = \frac{V_0^{EP}}{1 - \mu + V_0^{EP}},$$

equations (45), (46), and (47) characterize the Equal Pay equilibrium in the closed-form environment. Protected-worker welfare under Equal Pay is lower than under *laissez-faire* exactly when

$$p_0^{EP}(w^{EP} - b) < (p_0^{LF})^2 S_0. \tag{48}$$

Equivalently, this condition can be written as (39).

The disappearance of the disutility submarket reduces the set of active vacancies associated with disutility firms. Non-disutility firms may expand vacancy creation in response, but because vacancy creation is convex, this expansion is not infinitely elastic and need not fully offset the lost vacancies. Thus the Equal Pay equilibrium combines a wage-pooling effect with a congestion effect operating through p_N^{EP} and p_0^{EP}.

Quota Policy. Under the strict Quota Policy, the disutility protected-worker submarket is inactive, so

$$A_d^Q = V_d^Q = 0.$$

The effective quota condition is

$$p_N^Q = p_0^Q.$$

Let $p^Q = p_N^Q = p_0^Q$. Since $S_N = S_0$ and $\beta(\theta_N^Q) = \beta(\theta_0^Q) = p^Q$, one has

$$w_N^Q = w_0^Q = b + p^Q S_0.$$

Protected-worker expected utility is therefore

$$U_P^Q = b + (p^Q)^2 S_0 - \delta.$$

The quota reduces protected-worker welfare if and only if

$$(p^Q)^2 S_0 < (p_0^{LF})^2 S_0.$$

Since $S_0 > 0$, this is equivalent to (42). Thus, in the closed-form environment, the Quota Policy reduces protected-worker welfare exactly when the common job-finding probability generated by the constrained equilibrium is below the protected-worker *laissez-faire* probability in the non-disutility submarket.