WORKING PAPERS

N° 1798

Octobre 2026

“Equal Treatment, Unequal Family Duties: Gendered Screening and Selective Exclusion in the Workplace”

Francesca Barigozzi and Helmuth Cremer

Toulouse School of Economics

Equal Treatment, Unequal Family Duties: Gendered Screening and Selective Exclusion in the Workplace*

Francesca Barigozzi† Helmuth Cremer‡

October 24, 2026

Abstract

Can a ban on explicit gender discrimination in hiring and compensation unintentionally reinforce labor market inequalities when family duties remain unequally distributed across genders? This paper develops a theoretical model to address this question in a monopsonistic labor market where workers differ in gender-specific outside options and labor-supply costs—driven by family responsibilities—but exhibit the same underlying heterogeneity in effort costs. We compare outcomes under gender-based tagging and gender-neutral screening. When tagging is allowed, the firm offers efficient, gender-specific contracts. Under gender-neutral screening, by contrast, interacting incentive and participation constraints distort both the extensive and intensive margins of labor supply, amplifying gender gaps in employment, effort, labor income, and wages.

The model identifies five regimes and yields a paradox: as women's outside options approach men's, gender-neutral screening can intensify female exclusion. Improved female opportunities make the contracts intended for women more attractive to men. To prevent mimicking, the firm distorts these contracts and hires fewer women than under tagging. For a given efficiency type, the firm's cost of gender neutrality peaks when selective exclusion of that type begins, and gender gaps in employment and labor income widen. Unless family duties are shared more equally, equal treatment may therefore produce more unequal workplace outcomes.

Keywords: Heterogeneous family duties, converging outside options, gender neutrality, screening, exclusion.

***JEL*-Classification:** D82, J16, J42, J71, M54.

*We are grateful to participants in the seminar at the University of Bologna for their valuable comments and suggestions.

Financial support from the Chaire "*Marché des risques et creation de valeur*" of the FdR/SCOR is gratefully acknowledged. Helmuth Cremer also acknowledges the funding received by TSE from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program). Francesca Barigozzi also acknowledges financial support from the RFO funding scheme of the University of Bologna.

†University of Bologna, Italy, Email: francesca.barigozzi@unibo.it.

‡Toulouse School of Economics, France, Email: helmuth.cremer@tse-fr.eu.

1 Introduction

Labor markets remain deeply gendered despite decades of legal and societal evolutions. A vast empirical literature documents persistent asymmetries in both market and household behavior that shape the opportunity sets of men and women; see, among many others, Petrongolo (2019) and Bertrand (2020). Female job seekers tend to have lower reservation wages than male job seekers (Filippin and Ichino, 2005), a pattern consistent with the awareness of weaker labor-market outside options, driven by the gender wage gap and occupational segregation. Once employed, women supply fewer market hours than men, especially after motherhood. Evidence from Denmark, Sweden, Austria, and other countries shows that the "child penalty" translates into persistent reductions in female labor supply and earnings; see, among others, Kleven et al. (2019a), Kleven et al. (2019b), and Casarico and Lattanzio (2023). Moreover, income losses associated with caregiving may extend beyond childcare: later in life, women may also reduce labor supply or earnings to care for impaired household members, especially elderly parents; see Klimaviciute and Pestieau (2023) for a survey on long-term care. These labor-supply asymmetries are mirrored in wages: even among full-time employees, women earn substantially less than men, a robust empirical regularity summarized in Blau and Kahn (2017).

Household behavior mirrors and reinforces these labor-market inequalities. Across countries, time-use data consistently reveal that women devote significantly more hours than men to unpaid household work and informal childcare, even when they work full-time. In Italy, for example, full-time working mothers spend almost three more hours per day on unpaid work (household work plus childcare) than their partners; Barigozzi et al. (2023). Similar patterns hold in the United States, the United Kingdom, and Northern Europe; Aguiar and Hurst (2007), Giménez-Nadal and Sevilla (2012), Bianchi et al. (2012). These persistent asymmetries imply that women bear a higher marginal cost of supplying labor, since each hour worked crowds out time that must otherwise be allocated to household and caregiving duties.

Taken together, the empirical evidence suggests a characteristic feature of labor markets: women enter negotiation with lower outside options and face higher costs of supplying labor once employed. This combination provides the starting point for our theoretical analysis of how explicit anti-discrimination constraints shape optimal screening menus and worker sorting.

Modern anti-discrimination law restricts the use of gender as a contractual tag in employment relationships. In the European Union, Directive 2006/54/EC requires equal treatment of men and women in matters of employment and occupation and prohibits both direct and indirect discrimination on grounds of sex in access to employment, including "selection criteria and recruitment conditions," as well as in "employment and working conditions, including dismissals, as well as pay."[1] The same Directive also provides that job classification systems used to

[1]Directive 2006/54/EC, Art. 14(1), provides that "there shall be no direct or indirect discrimination on grounds of sex" in the public or private sectors in relation to, among other things, "conditions for access to employment, to

determine pay must rely on the same criteria for men and women and be designed so as to exclude sex-based discrimination. More recently, Directive (EU) 2023/970 on pay transparency has further operationalized this principle by requiring pay structures to allow comparisons on the basis of objective, gender-neutral criteria.[2]

We focus on the ban on explicit gender-contingent contracts. Gender is typically observable to the employer but, under gender-neutral contracting, it cannot be used as an explicit contractual tag: the employer cannot condition contracts explicitly on workers' gender. It can, however, offer a menu of contracts that is formally available to all workers and whose terms are defined in gender-neutral variables, such as wages and labor supply. For example, the employer can offer a menu containing both part-time and full-time contracts. In this sense, gender neutrality restricts the set of admissible contracts, but it does not eliminate self-selection and gendered sorting.

We document the effects of these mechanisms. If workers differ in non-contractible labor-supply costs that are systematically correlated with gender, a formally gender-neutral menu may induce gender-correlated choices. For instance, if a firm offers both part-time and full-time contracts, women—because of heavier family duties—may disproportionately select into part-time jobs. The resulting allocation is gender-neutral in its formal terms, but generates gendered outcomes through self-selection.

Our analysis should therefore be read as a positive analysis of the economic forces that can operate under formally gender-neutral contracting. We ask how the screening problem changes once explicit gender-contingent contracts are ruled out, and document that the remaining admissible instruments still induce sorting along gender-correlated constraints.

This form of gender-correlated self-selection is consistent with existing evidence. A 1999 Spanish law granting employment protection to workers who request a shorter workweek for family responsibilities generated a highly gendered take-up: almost exclusively women requested part-time hours. Fernández-Kranz and Rodríguez-Planas (2021) show that employers reacted by becoming 49% less likely to hire women of childbearing age, 40% more likely to separate from them, and 37% less likely to promote them to permanent contracts, increasing female non-employment relative to comparable men. This episode illustrates how formally gender-neutral or family-oriented policies may produce gendered labor-market effects when care burdens are unequally distributed. We propose a theoretical model to explain gendered outcomes generated by a ban on explicit gender-contingent contracts when systematic gender differences tend to

self-employment or to occupation, including selection criteria and recruitment conditions" and "employment and working conditions, including dismissals, as well as pay." Article 4 further states that, "where a job classification system is used for determining pay, it shall be based on the same criteria for both men and women and so drawn up as to exclude any discrimination on grounds of sex."

[2]Directive (EU) 2023/970, Art. 4(4), requires pay structures to be such as to enable an assessment of whether workers are in a comparable situation with regard to the value of work "on the basis of objective, gender-neutral criteria."

persist.

We take the employer's perspective. We study how anti-discrimination constraints shape the firm's optimal design of gender-neutral menus. The mechanism we emphasize does not rely on discriminatory preferences. It arises because gender cannot be used as a contractual tag, while some non-contractible determinants of labor supply—in particular family-duty costs—remain systematically correlated with gender.

We study this mechanism in a monopsonistic labor-market model in which male and female workers have the same underlying distribution of effort costs, but differ in outside options and in the additional cost of supplying labor associated with family responsibilities.[3]

We adopt a historical perspective motivated by the asymmetric evolution of market and household constraints. Women's labor-market opportunities have improved over time: participation has increased, and occupational opportunities have expanded. In the model, this is captured by an increase in women's outside options toward those of men. Within households, however, convergence has been slower: time-use evidence shows that women remain primarily responsible for unpaid household work and caregiving, including in dual-earner and full-time employed couples. Thus, market convergence has not been matched by an equally large reallocation of family responsibilities.

We study contracts along their three key components: the transfer (a nonlinear wage schedule), the instrument used to generate self-selection (the worker's effort, i.e., the number of hours worked), and the extensive margin, captured by the effort cost of the marginal worker. We first characterize the benchmark in which gender-based tagging is allowed. In that setting, the employer offers efficient, gender-specific contracts and extracts all worker surplus. Efficient contracts are such that male workers earn more and exert more effort (i.e., work longer hours) than their female counterparts, who suffer from heavier family duties. The efficient marginal worker is defined as the least productive agent for whom employment is still profitable for the firm, given the worker's outside option and the cost of supplying labor. Intuitively, men's labor supply is cheaper because they have lower family-duty costs, but women are easier to attract because they have lower outside options. Hence, the firm hires more women than men only when the outside-option advantage dominates the family-duty cost disadvantage. Moreover, an increase in women's outside option reduces the benchmark female-employment cutoff because women become more costly to recruit.

Under gender neutrality, tagging is prohibited, and the employer must offer the same menu of contracts to both genders. This generates a rich set of incentive and participation constraints whose interaction depends on how close female outside options are to those of male workers. Our analysis reveals five possible regimes, each defined by a different set of binding incentive and participation constraints. In one regime (B), incentive compatibility constraints are slack, and the benchmark efficient contracts remain implementable despite gender neutrality. In the

[3]We abstract from efficiency screening to isolate screening on non-contractible family-duty costs.

remaining regimes ($R1$, $R2$, $R4$, and $R5$), however, incentive constraints bind and generate distortions not only in effort and wages but—crucially—along the extensive margin. Relative to standard screening models with type-dependent participation constraints, our contribution is to combine observable worker efficiency, gender-correlated effort costs, a legally motivated restriction on explicit gender tagging, and endogenous gender-specific employment cutoffs. Together, these elements generate a new extensive-margin distortion: selective exclusion relative to the tagging benchmark. Because worker efficiency is observable, the firm's problem separates across efficiency levels, which yields a simple hiring rule: at each efficiency level, the firm employs women only if the profit they generate exceeds the cost they impose through men's incentive constraint, namely the distortion of women's hours and any rent paid to men. As female outside options increase and approach those of men, the model progressively transitions from Regime R5 to Regime R1.

To understand our results, recall that men's effort is cheaper because they bear fewer family duties, whereas women are easier to attract because they have lower outside options. Hence, when the firm hires both male and female workers, wages in the screening contracts depend on both the costs of effort and the outside options.

When the gender gap in outside options is large, as was historically the case, Regimes $R4$ and especially $R5$ may arise. In both regimes, female workers have an incentive to choose the contracts designed for men, since these contracts offer higher wages due to men's lower effort costs and higher outside options. To deter such mimicking, the firm raises the effort required from men — a countervailing-incentives effect — and restricts male employment, generating additional male exclusion relative to the tagging benchmark.[4]

In contemporary labor markets, the more empirically relevant cases are Regimes $R2$ and $R1$, which arise when women's outside options are close to those of men. In these regimes, male workers prefer the contracts designed for women, which require less effort — that is, fewer working hours — while still offering a relatively high wage, because women's outside options are high. The firm responds by reducing female effort and lowering the female-employment cutoff below its tagging benchmark, thereby limiting the cost of satisfying men's incentive constraint. Although a higher female outside option reduces female employment even under tagging, the cutoff falls proportionally faster under gender neutrality, generating additional female exclusion: a higher female outside option raises not only the cost of hiring women, but also, through men's incentive constraint, the cost of deterring male mimicking. Thus, paradoxically, improvements in women's outside options magnify selective female exclusion, with the firm's profit loss relative to tagging peaking when exclusion begins.

[4]Regime $R5$, however, is unlikely to be empirically plausible, as its existence requires a large gender gap in outside options combined with a relatively small gender gap in family-duty costs. One would instead expect these two gaps to be positively correlated: larger inequalities in the labor market are likely to be associated with larger inequalities within the household.

We also show that gender-neutral contracting can increase *wage inequality*. Under tagging, efficient effort implies that the gender wage gap decreases monotonically as women's outside options improve. Under gender neutrality, incentive compatibility distorts male effort upward when women are the mimickers and female effort downward when men are the mimickers. Because wages compensate effort costs and reflect outside options and information rents, the gender wage gap widens even as outside options converge.

Wage and utility gaps need not move together. Improving women's outside options tends to narrow the gap in *indirect utility*, but employed women continue to earn less because utility convergence is driven by lower effort costs (lower labor supply). Moreover, higher female outside options intensify selective female exclusion, so some women cease to be employed before their within-employment utility converges to men's. Thus, narrower utility gaps can coexist with persistent wage inequality and widening employment gaps.

Overall, identical contractual treatment need not produce equal outcomes when non-contractible constraints remain gendered. Gender-neutral workplace policies should therefore be accompanied by policies that promote a more equal sharing of family responsibilities. Only when family duties are shared more symmetrically can gender-neutral hiring and compensation policies deliver on their promise.

1.1 Related literature

Our paper first relates to the contract-theory literature on type-dependent participation constraints and countervailing incentives. When outside options vary across types, the conventional ranking of information rents may be reversed, and different participation and incentive constraints may bind across parameter regions (Lewis and Sappington, 1989; Jullien, 2000; Laffont and Martimort, 2002). We retain this logic but introduce an observable efficiency dimension and endogenous type-specific exclusion cutoffs. Our tagging benchmark follows the logic of conditioning allocations on observable group characteristics (Akerlof, 1978; Cremer et al., 2010). In a related application, Alesina et al. (2011) show that gender-based taxation can be optimal and can induce a more balanced division of market and household work by affecting intrafamily bargaining. Whereas they study how gender-specific taxes reshape household choices, we examine how prohibiting gender-contingent employment contracts affects screening when family duties remain unequally distributed. The resulting distortions operate not only through effort and information rents but also through the gender composition of employment.

Second, our monopsonistic setting connects the paper to work on employer market power and gender wage differentials (Manning, 2003; Barth and Dale-Olsen, 2009). Whereas this literature emphasizes gender differences in labor-supply elasticities and mobility frictions, we hold the distribution of worker efficiency fixed across genders and use monopsony as the environment in which the employer designs a screening menu. Wage and employment gaps arise from the interaction of gender-correlated family-duty costs, gendered outside options, and a ban on

gender-contingent contracts.

Third, the paper relates to evidence that women value flexibility and sort into alternative work arrangements, often accepting lower earnings in return (Mas and Pallais, 2017; Wiswall and Zafar, 2018; Flabbi and Moro, 2012). Theoretical work shows how household responsibilities and employers' expectations can generate gender gaps in hiring and promotion (Albanesi and Olivetti, 2009; Lommerud et al., 2015). Closest to our analysis, Frech and Maideu-Morera (2024) study dynamic contracting when stochastic and unobservable flexibility needs differ by gender, showing that contracts not tailored to women's needs can generate wage divergence and employment termination. Our mechanism is static and focuses instead on how changing outside options and a non-tagging constraint reshape a common screening menu and its endogenous employment cutoffs.

Finally, our results provide a contractual mechanism for the unequal effects of formally gender-neutral policies documented empirically. Antecol et al. (2018) show that gender-neutral tenure-clock extensions increased male but reduced female tenure rates, while Fernández-Kranz and Rodríguez-Planas (2021) document adverse hiring, separation, and promotion effects associated with a family-oriented right to reduced working hours. We show how gender gaps in labor market outcomes can arise without discriminatory preferences: when family constraints remain gendered, a common menu can induce gendered self-selection and additional exclusion relative to gender tagging.

2 The model

We consider a sector in which a single firm employs both female and male workers.

2.1 Workers

The utility of female workers is given by

$$U_f = w - \frac{1}{2}\theta e^2 - \gamma_f e, \tag{1}$$

while that of male workers is

$$U_m = w - \frac{1}{2}\theta e^2 - \gamma_m e, \tag{2}$$

where w denotes the wage and e denotes effort, which can also be interpreted as labor supply. Supplying effort is costly. The second and third terms in equations (1) and (2) capture the disutility associated with effort provision.

There is an equal unit mass of male and female workers. Conditional on gender, workers are heterogeneous with respect to the cost-of-effort parameter θ, which is uniformly distributed over the common support $[\underline{\theta}, \overline{\theta}]$, where $0 < \underline{\theta} < \overline{\theta}$.

The parameter θ can be interpreted as an inverse measure of worker efficiency. Throughout the paper, we assume that θ is observable. This assumption allows us to abstract from

standard efficiency screening and focus on the screening problem generated by gender-correlated family-duty costs. The parameter can be interpreted as a characteristic observed by the employer through education, experience, occupation, or job-specific performance indicators. This assumption also makes the extensive margin transparent: for each level of effort cost, the firm decides whether employment is profitable under the relevant incentive and participation constraints.

The third term in the utility function is gender-specific and captures the additional cost of labor supply associated with family duties. It is governed by the parameter γ, with

$$1 > \gamma_f \geq \gamma_m \geq 0.$$

Thus, for female workers, the cost of any given level of labor supply is weakly higher, reflecting the fact that labor supply may come at the expense of family duties, which tend to be greater for women (Aguiar and Hurst, 2007; Gimenez-Nadal and Sevilla, 2012; Bianchi et al., 2012).[5]

Outside options are also gender-specific. We assume that

$$\overline{U}_m \equiv \overline{U}_m^k \geq \overline{U}_f \geq \overline{U}_f^o > 0,$$

so that employment opportunities in the outside-option sector are better for men than for women (among many others, Blau and Kahn, 2017; Bertrand, 2020).

Historically, the labor market has evolved from a context characterized by substantial gender differences in outside options toward one with more equal opportunities for men and women; see Goldin (2014). To capture this trend, we keep the male outside option, $\overline{U}_m \equiv \overline{U}_m^k$, fixed, while allowing the female outside option, $\overline{U}_f$, to increase over time. In the past, $\overline{U}_f$ was low, with $\overline{U}_f = \overline{U}_f^o$, implying substantial heterogeneity in outside options: $\overline{U}_m^k - \overline{U}_f^o \gg 0$.

More recently, the economy has moved toward a situation in which this heterogeneity gradually vanishes, as $\overline{U}_f$ approaches $\overline{U}_m^k$. Hence, outside options are characterized by

$$\overline{U}_m \equiv \overline{U}_m^k, \qquad \overline{U}_f \in \left[\overline{U}_f^o, \overline{U}_m^k\right], \qquad \overline{U}_f^o > 0.$$

2.2 The employer

Workers' effort, or labor supply, is the only input used in production. Effort is contractible. The firm's production function exhibits constant returns to effort, so that output is given by

$$q(e) = e,$$

with the unit value of output normalized to one.

The profit generated by a worker exerting effort e is therefore

$$\pi(e) = e - w(e),$$

[5]The gender-specific parameters γ_f and γ_m are reduced-form group averages capturing systematic differences in family constraints. For tractability, the model abstracts from within-gender heterogeneity and assumes a deterministic mapping between gender and γ. This simplification isolates the screening mechanism and does not imply that every woman faces higher family-duty costs than every man.

where the wage depends on the level of effort. The monopsonist maximizes total profits, given by the integral of individual worker profits over the set of employed workers.

We consider a monopsonistic labor market in order to focus sharply on the firm's contract-design problem. The monopsonist chooses the optimal screening menu and therefore internalizes the key trade-offs generated by gender neutrality: effort distortions, information rents, and the extensive margin of employment. This setting provides a transparent environment in which to isolate the mechanism linking gender-neutral contracting, gendered family-duty costs, and selective exclusion. The role of employer competition is discussed in the concluding remarks.

As a benchmark, we first consider a setting in which the firm can use gender as a tag, meaning that the monopsonist can condition the contract offered to a worker on the worker's gender.

3 Monopsony with tagging (Benchmark)

Suppose first that gender-based tagging is feasible. Since both workers' type θ and gender are observable, and no gender-neutrality constraint is imposed, the firm can condition contracts on gender.

An allocation consists of two triplets, $\{e_g(\theta), w_g(\theta), \widetilde{\theta}_g\}_{g\in\{f,m\}}$, where $e_g(\theta)$ and $w_g(\theta)$ denote, respectively, labor supply (effort) and compensation for a worker of gender g and type θ, while $\widetilde{\theta}_g$ is the gender-specific marginal worker. Workers of gender g with $\theta \leq \widetilde{\theta}_g$, namely those who are more productive than the marginal worker, are employed by the firm; no contract is offered to the remaining workers.

Under full information and in the absence of gender neutrality (GN), the firm can extract the entire surplus. Profit maximization, therefore, yields the efficient allocation in terms of both effort levels and marginal workers. Compensation is then set so that each employed worker obtains exactly her or his reservation utility.

The monopsonist's program is:[6]

$$\max_{e_g(\theta),w_g(\theta),\widetilde{\theta}_g} \Pi = \int_{\underline{\theta}}^{\widetilde{\theta}_f} (e_f(\theta) - w_f(\theta))d\theta + \int_{\underline{\theta}}^{\widetilde{\theta}_m} (e_m(\theta) - w_m(\theta))d\theta, \tag{3}$$

$$\text{s.t.}: \quad w_f(\theta) - \frac{1}{2}\theta\left(e_f(\theta)\right)^2 - \gamma_f e_f(\theta) \geq \overline{U}_f \quad \forall\theta, \quad PC_f(\theta)$$

$$w_m(\theta) - \frac{1}{2}\theta\left(e_m(\theta)\right)^2 - \gamma_m e_m(\theta) \geq \overline{U}_m \quad \forall\theta, \quad PC_m(\theta)$$

From the binding participation constraints we obtain the wage levels

$$w_f(\theta) = \frac{1}{2}\theta\left(e_f\left(\theta\right)\right)^2 + \gamma_f e_f\left(\theta\right) + \overline{U}_f, \tag{4}$$

$$w_m(\theta) = \frac{1}{2}\theta\left(e_m\left(\theta\right)\right)^2 + \gamma_m e_m\left(\theta\right) + \overline{U}_m. \tag{5}$$

[6]Because the density is identical and constant for both genders, it multiplies the firm's objective by a common positive factor and is omitted from the profit expressions.

Substituting (4) and (5) into the monopsonist's program (3), the first-order conditions (FOCs) with respect to labor supplies imply:

$$e_g^B(\theta) = \frac{1-\gamma_g}{\theta}, \quad g \in \{f, m\}. \tag{6}$$

Note that labor supply is decreasing in θ and that a female worker of a given effort-cost level provides a weakly lower labor supply than a male worker of the same θ:

$$e_f^B(\theta) \leq e_m^B(\theta) \qquad \forall \theta.$$

To understand this result, recall that all employed workers receive their respective outside options. They must therefore be compensated for the full cost of labor supply, including the term $\gamma_g e$. When $\gamma_f > \gamma_m$, the marginal cost of labor supply is higher for women, and their efficient labor supply is strictly lower.

Plugging the efficient labor supplies defined in (6) into equations (4) and (5) yields the efficient wages:

$$w_g^B(\theta) = \frac{1}{2}\frac{(1-\gamma_g)^2}{\theta} + \gamma_g \frac{1-\gamma_g}{\theta} + \overline{U}_g, \quad g \in \{f, m\}. \tag{7}$$

Using (7), the difference between male and female wages can be written as

$$w_m^B(\theta) - w_f^B(\theta) = \overline{U}_m^k - \overline{U}_f + \frac{\gamma_f^2 - \gamma_m^2}{2\theta} \geq 0.$$

Thus, male workers receive a weakly higher wage than female workers. The inequality is strict whenever $\overline{U}_f < \overline{U}_m^k$ or $\gamma_f > \gamma_m$; wages are equal only in the degenerate case $\overline{U}_f = \overline{U}_m^k$ and $\gamma_f = \gamma_m$.

Substituting the efficient wages and effort levels into the monopsonist's profits and rearranging, one can express profits as a function of the two marginal workers $\widetilde{\theta}_f$ and $\widetilde{\theta}_m$:

$$\Pi = \int_{\underline{\theta}}^{\widetilde{\theta}_f} \left(\frac{1}{2}\frac{(1-\gamma_f)^2}{\theta} - \overline{U}_f \right) d\theta + \int_{\underline{\theta}}^{\widetilde{\theta}_m} \left(\frac{1}{2}\frac{(1-\gamma_m)^2}{\theta} - \overline{U}_m \right) d\theta.$$

Differentiating with respect to $\widetilde{\theta}_f$ and $\widetilde{\theta}_m$, and solving by assuming an interior solution for both genders yields:

$$\widetilde{\theta}_f^B = \frac{1}{2}\frac{(1-\gamma_f)^2}{\overline{U}_f}; \tag{8}$$

$$\widetilde{\theta}_m^B = \frac{1}{2}\frac{(1-\gamma_m)^2}{\overline{U}_m}. \tag{9}$$

An interior solution such that $\underline{\theta} \leq \widetilde{\theta}_f^B, \widetilde{\theta}_m^B \leq \overline{\theta}$ arises if the outside option is sufficiently high to prevent the firm from hiring all workers, yet not so high as to deter hiring altogether.

The results obtained so far are summarized in the following proposition.

Proposition 1 ***Benchmark contracts.*** *When tagging is possible, gender-specific contracts* $\left\{w_g^B(\theta), e_g^B(\theta), \widetilde{\theta}_g^B\right\}_{g\in\{f,m\}}$ *are offered to workers with* $\theta \leq \widetilde{\theta}_g^B$, $g \in \{f,m\}$, *and are such that the monopsonist extracts all the workers' surplus.*

(i) Labor supplies are efficient *and given by expression (6). Men supply weakly more labor than women, with strict inequality when* $\gamma_f > \gamma_m$.

(ii) All workers receive their respective reservation utilities, and men's wages are weakly higher than women's. The wage inequality is strict whenever $\overline{U}_f < \overline{U}_m^k$ *or* $\gamma_f > \gamma_m$.

(iii) Marginal workers are given by expressions (8) and (9). The monopsonist hires more (less) women than men when the heterogeneity in outside options is larger (lower) than the heterogeneity in the additional cost of effort due to family duty:

$$\widetilde{\theta}_f^B \gtrless \widetilde{\theta}_m^B \Longleftrightarrow \frac{\overline{U}_m}{\overline{U}_f} \gtrless \frac{(1-\gamma_m)^2}{(1-\gamma_f)^2}.$$

To understand point *(iii)*, consider first the case where $\overline{U}_m = \overline{U}_f$. In this case, men and women have the same outside option, but men's labor cost is lower. As a result, the firm prefers to hire more men than women. However, as $\overline{U}_f$ decreases relative to $\overline{U}_m$, hiring women becomes more attractive. When $\overline{U}_f$ is sufficiently lower than $\overline{U}_m$, the effect of heterogeneity in outside options outweighs that of labor cost differences, and hiring more women than men becomes optimal.

Note that, from (8), an improvement in women's outside option reduces female employment because satisfying women's participation constraint becomes too costly. As we will see, GN may generate an additional employment effect by lowering the female marginal-worker cutoff below this benchmark level. In the analysis below, we distinguish this screening-induced distortion from the reduction in employment that already occurs under tagging.

4 Monopsony with gender neutrality

We now assume that the firm cannot condition contracts on gender. We solve the firm's problem in three steps, after a preliminary observation.

As a preliminary observation, we show that the problem separates across effort-cost levels: because θ is observable and every constraint refers to a single value of θ, the firm can solve a separate, small problem for each θ (Section 4.1). *Step 1* shows that a gender that the firm chooses not to employ never wants the contract offered to the other gender (Section 4.2). *Step 2* characterizes the best menu of contracts when both genders are employed; depending on θ and $\overline{U}_f$, this menu takes one of five forms, which we call *regimes* (Section 4.3). *Step 3* compares the profit from employing both genders with the profit from employing only one of them, which determines who is hired, that is, the extensive margin (Section 4.4).

Throughout this section, we assume that $\gamma_f > \gamma_m$,[7] and we use the notation

$$\Delta\gamma \equiv \gamma_f - \gamma_m > 0, \qquad \Delta\overline{U} \equiv \overline{U}_m^k - \overline{U}_f \geq 0.$$

Effort is constrained to be non-negative. All cutoff formulas below describe the latent value of θ at which the relevant profit becomes zero. If a latent cutoff lies outside the support $[\underline{\theta}, \overline{\theta}]$, the effective cutoff is obtained by truncating it at the appropriate boundary of the support.[8]

4.1 The problem at each effort-cost level

Under gender neutrality (GN), the firm must offer the same menu of contracts to men and women. Although gender is, in principle, observable, this constraint makes the problem formally equivalent to one in which gender is not observed. The employer therefore effectively screens on γ_g, while θ remains observable. Because θ is observable and can legally be used, the firm can offer different contracts to workers with different values of θ. For each θ the firm offers a single menu of contracts, available to all workers of that type regardless of gender. By the Revelation Principle, this menu contains at most two contracts, one intended for each gender, and a single contract when only one gender is employed. Offering more contracts cannot help the firm: any additional contract only adds an option that some worker might prefer to the contract intended for her or him.

Let

$$U_g(\theta) \equiv w_g(\theta) - \frac{1}{2}\theta \left(e_g(\theta)\right)^2 - \gamma_g e_g(\theta), \qquad g \in \{f, m\},$$

denote the utility that a worker of gender g and type θ obtains from the contract intended for her or him. At any θ at which both genders are employed, the two participation constraints are

$$U_f(\theta) \geq \overline{U}_f, \qquad U_m(\theta) \geq \overline{U}_m^k,$$

and the incentive constraints can be written as

$$IC_f(\theta): \;\; U_f(\theta) \geq U_m(\theta) - \Delta\gamma\, e_m(\theta), \qquad\qquad IC_m(\theta): \;\; U_m(\theta) \geq U_f(\theta) + \Delta\gamma\, e_f(\theta). \tag{10}$$

A woman who takes the contract designed for men obtains the men's utility minus $\Delta\gamma$ for each unit of effort, because each hour of work is more costly for her. Conversely, a man who takes the contract designed for women obtains the women's utility plus $\Delta\gamma$ for each unit of effort. Hence, the contract designed for women is attractive to men when it requires many hours, while the contract designed for men is attractive to women when the utility gap between the two contracts outweighs the extra family-duty cost that women would bear on the long male hours.

[7] If $\gamma_f = \gamma_m$, men and women differ only in their outside options. The incentive constraints then require that both genders obtain the same utility, so that gender neutrality simply forces the firm to leave women the same utility as men.

[8] Strict comparisons of employment levels apply when the relevant cutoffs lie inside the support of θ. If a cutoff lies outside the support, employment may be zero or equal to the full mass. Thus, different cutoffs may correspond to the same employment level.

When only one gender is employed at a given θ, only one contract is offered at that θ. Because this contract is formally available to workers of either gender, the worker who is not employed must prefer her or his outside option to it. If only men are employed, the following *no-entry* constraint for women must therefore hold:

$$U_m(\theta) - \Delta\gamma\, e_m(\theta) \leq \overline{U}_f. \tag{NE$_f$}$$

Symmetrically, if only women are employed, men must prefer their outside option:

$$U_f(\theta) + \Delta\gamma\, e_f(\theta) \leq \overline{U}^k_m. \tag{NE$_m$}$$

These no-entry constraints ensure that employment differences arise through voluntary self-selection from a gender-neutral menu.

The problem separates across effort-cost levels. Neither the objective nor the constraints link workers with different values of θ: a worker of type θ can choose only among the contracts offered to that type. The firm can therefore solve a separate problem for each θ and then integrate the resulting pointwise profit. Importantly, we do not impose cutoff employment sets at this stage. The cutoff structure will follow from the monotonicity results in Lemma 3.

To describe the problem at a given θ, let

$$S_g(e,\theta) \equiv (1-\gamma_g)e - \frac{1}{2}\theta e^2$$

denote the surplus generated by a worker of gender g who exerts effort e (output minus the worker's effort and family-duty costs). The profit earned on a worker of gender g is $e_g - w_g = S_g(e_g,\theta) - U_g$: surplus minus the utility left to the worker. Under tagging, the firm sets effort at its efficient level $e^B_g(\theta)$ and leaves the worker her or his outside option, so that the profit per worker is

$$\pi^B_g(\theta) \equiv \max_e S_g(e,\theta) - \overline{U}_g = \frac{(1-\gamma_g)^2}{2\theta} - \overline{U}_g, \qquad g \in \{f,m\}, \tag{11}$$

which is positive if and only if $\theta < \widetilde{\theta}^B_g$. Because the surplus is quadratic in effort, a distortion of effort away from its efficient level reduces surplus by

$$S_g(e^B_g(\theta),\theta) - S_g(e,\theta) = \frac{1}{2}\theta\left(e - e^B_g(\theta)\right)^2. \tag{12}$$

At each θ, the firm has four options:

(a) *Employ both genders* with two self-selecting contracts. The best profit from this option is

$$\begin{aligned} V(\theta) \equiv \max_{e_f,e_m,U_f,U_m} \; & [S_f(e_f,\theta) - U_f + S_m(e_m,\theta) - U_m] \\ \text{s.t.} \;\; & PC_f(\theta),\, PC_m(\theta),\, IC_f(\theta),\, IC_m(\theta), \\ & e_f \geq 0, \quad e_m \geq 0. \end{aligned} \tag{13}$$

(b) *Employ only men.* Let $V_M(\theta)$ denote the maximum profit from a single contract (e_m, U_m) satisfying $PC_m(\theta)$, $NE_f(\theta)$, and $e_m \geq 0$.

(c) *Employ only women.* Let $V_F(\theta)$ denote the maximum profit from a single contract (e_f, U_f) satisfying $PC_f(\theta)$, $NE_m(\theta)$, and $e_f \geq 0$.

(d) *Employ nobody*, which yields zero profit.

Before solving the single-gender problems, note that, without presupposing any cutoff structure, total profit under gender neutrality can therefore be written as

$$\Pi^{GN} = \int_{\underline{\theta}}^{\overline{\theta}} \max\{V(\theta), V_M(\theta), V_F(\theta), 0\}\, d\theta. \tag{14}$$

Lemma 1 will show that the two single-gender values can be replaced in this pointwise comparison by the corresponding benchmark profits. The employment cutoffs will then emerge endogenously from the values of θ at which the maximizing option changes.

Problem (13) is a standard two-type screening problem with type-dependent outside options, in which the "type" is given by γ, which captures the family-duty-related cost of effort and is fully determined by gender. Laffont and Martimort (2002, Section 3.3) show that, in such a problem, the agents' outside options can lead to five distinct regimes. These range from a standard case, where the incentive constraint of the "low-cost" type and the participation constraint of the "high-cost" type bind, to a countervailing-incentives regime, where the reverse holds. Between these extremes lie three additional cases: one in which the full-information optimum is incentive compatible, and two transitional regimes in which both participation constraints and one incentive constraint bind. Our setting adds two elements to this structure. First, the regime depends on θ as well as on $\overline{U}_f$. Second, the firm also decides whom to employ at each θ, which generates the extensive margin.

4.2 Step 1: Exclusion does not require contract distortions

Options (b) and (c) involve a no-entry constraint and might therefore require the firm to distort the contract of the employed gender in order to keep the other gender out. The following lemma shows that such a distortion is never optimal. The proof uses only the definition of $V(\theta)$ in (13), not its explicit solution, which we derive in Step 2.

Lemma 1 ***No distortion under exclusion.*** *Fix θ. (i) If $\pi_m^B(\theta) > 0$ and women strictly prefer the benchmark male contract to their outside option, then*

$$V(\theta) \geq 2\pi_m^B(\theta) > \pi_m^B(\theta),$$

so employing only men cannot be optimal. The symmetric result holds if men strictly prefer the benchmark female contract.

(ii) The firm's maximum profit at θ is

$$\Pi^*(\theta) = \max\left\{V(\theta),\, \pi_m^B(\theta),\, \pi_f^B(\theta),\, 0\right\}. \tag{15}$$

Whenever employing only one gender is strictly optimal, the employed gender receives its benchmark contract, and the other gender weakly prefers its outside option to that contract.

To see the argument, suppose that women strictly prefer the benchmark male contract. The firm can offer this single contract to both genders. Both accept it, and each generates profit $\pi_m^B(\theta)$, because profit from a given effort–wage pair does not depend on the worker's gender. Hence the firm obtains $2\pi_m^B(\theta)$. By contrast, any contract that employs only men generates at most $\pi_m^B(\theta)$, since the benchmark contract already maximizes profit from men. Employing only men therefore cannot be optimal. Equivalently, if employing only men is optimal, the benchmark male contract must already satisfy women's no-entry constraint:

$$U_f\big(w_m^B(\theta), e_m^B(\theta)\big) \leq \overline{U}_f.$$

The argument for employing only women is symmetric.

Equation (15) therefore reduces the exclusion decision to comparing the profit from employing both genders, the two single-gender benchmark profits, and zero. It remains to compute $V(\theta)$.

4.3 Step 2: The optimal menu when both genders are employed

A key ingredient in the analysis is the set of conditions that a menu must satisfy to be incentive compatible. These conditions are provided in Lemma 2, established in Appendix B.

Lemma 2 ***Incentive compatibility.*** *The two incentive constraints are satisfied if and only if*

$$U_m(\theta) - \Delta\gamma e_m(\theta) \leq U_f(\theta) \leq U_m(\theta) - \Delta\gamma e_f(\theta); \tag{16}$$

Since effort is non-negative, this implies $U_m(\theta) \geq U_f(\theta)$; since $\Delta\gamma > 0$, it also implies $e_m(\theta) \geq e_f(\theta)$.

The second implication means that any gender-neutral menu sorts workers by hours: the contract chosen by men requires at least as much effort as the contract chosen by women.

When is the tagging allocation still feasible? We first ask whether the benchmark contracts of Section 3 satisfy the incentive constraints. Within the region where they do, the solution with gender tagging remains implementable under the GN constraint and continues to solve the firm's problem. The relevant conditions are stated in the following proposition, which is proved in Appendix C. Figure 1 provides a graphical illustration.

Proposition 2 ***Incentive compatible benchmark contracts.*** *(i) Benchmark contracts with* $e_g(\theta) = e_g^B(\theta)$ *and* $U_g(\theta) = \overline{U}_g$, $g = f, m$, *are incentive compatible for:*

$$\overline{U}_m^k - \frac{\Delta\gamma(1-\gamma_m)}{\theta} \leq \overline{U}_f \leq \overline{U}_m^k - \frac{\Delta\gamma(1-\gamma_f)}{\theta}. \tag{17}$$

The length of the interval such that Benchmark contracts are incentive compatible is $(\gamma_f - \gamma_m)^2/\theta$, *which decreases with* θ. *When the first inequality in (17) is violated,* $IC_f(\theta)$ *is violated by the benchmark contracts; when the second inequality is violated,* $IC_m(\theta)$ *is violated.*
(ii) In the θ *dimension, condition (17) can be expressed as*

$$\theta^B_{IC_m}\left(\overline{U}_f\right) \equiv \frac{(\gamma_f - \gamma_m)(1-\gamma_f)}{\overline{U}_m^k - \overline{U}_f} \leq \theta \leq \theta^B_{IC_f}\left(\overline{U}_f\right) \equiv \frac{(\gamma_f - \gamma_m)(1-\gamma_m)}{\overline{U}_m^k - \overline{U}_f}, \tag{18}$$

where the functions $\theta^B_{IC_m}\left(\overline{U}_f\right)$ *and* $\theta^B_{IC_f}\left(\overline{U}_f\right)$ *are increasing and convex.*
(iii) The two benchmark employment cutoffs $\widetilde{\theta}_f^B$ *and* $\widetilde{\theta}_m^B$ *defined in (8) and (9) intersect at point* $I = (\overline{U}_f^I, \theta^I)$ *of Figure 1, where* $\tilde{\theta}_f^B(\overline{U}_f^I) = \tilde{\theta}_m^B = \theta^I$ *and* $\theta^B_{IC_m}(\overline{U}_f^I) < \theta^I < \theta^B_{IC_f}(\overline{U}_f^I)$.
Thus, the benchmark cutoffs intersect strictly inside the region in which the benchmark contracts are incentive compatible.

The logic of Proposition 2 explains why the mimicking gender depends on θ. Under the benchmark contracts, a man who takes the contract designed for women obtains $\overline{U}_f + \Delta\gamma e_f^B(\theta)$: he receives the women's outside option plus a gain of $\Delta\gamma$ on each hour, because his family-duty cost is lower. He prefers this contract to his own when $\Delta\gamma e_f^B(\theta) > \Delta\overline{U}$, that is, when $\theta < \theta^B_{IC_m}(\overline{U}_f)$. His gain is proportional to women's hours, which are long when workers are efficient (low θ). Men are therefore tempted by the female contract at low values of θ, and more so when $\overline{U}_f$ is close to $\overline{U}_m^k$. Conversely, a woman who takes the contract designed for men obtains $\overline{U}_m^k - \Delta\gamma e_m^B(\theta)$: she gains the outside-option gap but bears a higher family-duty cost on each of the long male hours. She prefers this contract when $\Delta\overline{U} > \Delta\gamma e_m^B(\theta)$, that is, when $\theta > \theta^B_{IC_f}(\overline{U}_f)$. Women are therefore tempted by the male contract at high values of θ, where hours are short, and more so when their outside option is low. In the region between the two frontiers, which contains the gray area of Figure 1, neither gender wants the other's contract and the benchmark allocation is implemented.

Figure 1 combines the two incentive-compatibility frontiers with the benchmark employment cutoffs, $\widetilde{\theta}_f^B(\overline{U}_f)$ and $\widetilde{\theta}_m^B$. The latter is a horizontal line because it does not depend on $\overline{U}_f$. A horizontal movement holds worker efficiency fixed and increases women's outside option, whereas a vertical movement holds outside options fixed and compares workers with different effort costs. The labeled points in the figure are intersections among the two benchmark cutoffs and the two incentive-compatibility frontiers. Points A and C are the intersections of the male cutoff $\tilde{\theta}_m^B$ with $\theta^B_{IC_f}(\overline{U}_f)$ and $\theta^B_{IC_m}(\overline{U}_f)$, respectively. Points H and K are the corresponding intersections of the female cutoff $\tilde{\theta}_f^B(\overline{U}_f)$ with these two frontiers, while point I is the intersection of the two benchmark cutoffs. Proposition 2(iii) gives the coordinates of I, and Appendix C derives the

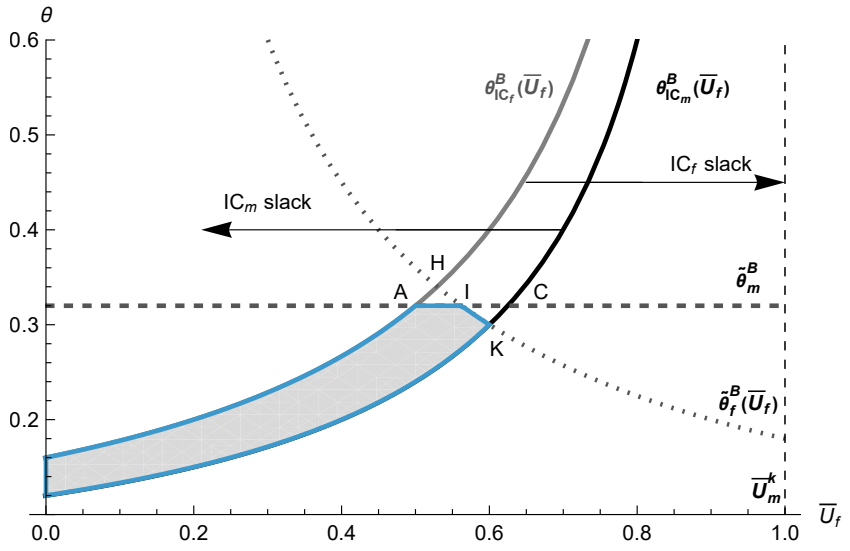


Figure 1: **Incentive compatible Benchmark contracts in the space $(\overline{U}_f, \theta)$.** The figure depicts the functions describing the efficient marginal workers, $\widetilde{\theta}^B_g$, $g \in \{f, m\}$ (see equations (8) and (9)), together with the two frontiers where the incentive constraints are binding $\theta^B_{IC_g}(\overline{U}_f)$, $g \in \{f, m\}$ (see equation (18)). In the gray area, both genders are employed, and Benchmark contracts are incentive compatible. The small area delimited by AIH is the region where only women are employed, while the small area delimited by ICK is the region where only men are employed. In both areas AIH and ICK, and in the gray region, workers receive Benchmark contracts. Above HIC, neither gender is hired. In the figure, $\overline{U}^k_m = 1$, $\gamma_f = 0.4$, and $\gamma_m = 0.2$.

other coordinates and establishes the ordering $\overline{U}^A_f < \overline{U}^H_f < \overline{U}^I_f < \overline{U}^K_f < \overline{U}^C_f$. In particular, point I lies inside the incentive-compatible Benchmark region: the two benchmark employment cutoffs cross while the benchmark contracts remain implementable under gender neutrality.

Recall that gender g is profitably employed under tagging if and only if $\theta < \widetilde{\theta}^B_g$. Hence, within the region where the benchmark contracts are incentive compatible, both genders are employed when θ lies below both cutoffs. In region AIH,

$$\widetilde{\theta}^B_m \leq \theta < \widetilde{\theta}^B_f(\overline{U}_f),$$

so only women are profitably employed. In region ICK,

$$\widetilde{\theta}^B_f(\overline{U}_f) \leq \theta < \widetilde{\theta}^B_m,$$

so only men are profitably employed. From Lemma 1, in both regions, the relevant no-entry constraint is satisfied, and the employed gender receives its benchmark contract without distortion.

The five regimes. Outside the region defined by (18), the firm must modify the menu. Two further frontiers determine how it does so. As shown in Proposition 3 established in Appendix D, these frontiers are obtained by setting the information rent of the mimicking gender equal to

zero, so that its participation constraint just binds:

$$\theta_{PC_m}(\overline{U}_f) \equiv \Delta\gamma \frac{1-2\gamma_f+\gamma_m}{\overline{U}_m^k - \overline{U}_f}, \qquad \theta_{PC_f}(\overline{U}_f) \equiv \Delta\gamma \frac{1-2\gamma_m+\gamma_f}{\overline{U}_m^k - \overline{U}_f}. \tag{19}$$

When $\overline{U}_f < \overline{U}_m^k$, the four frontiers, depicted together in Figure 2, are ordered as follows:

$$\theta_{PC_m}(\overline{U}_f) < \theta^B_{IC_m}(\overline{U}_f) < \theta^B_{IC_f}(\overline{U}_f) < \theta_{PC_f}(\overline{U}_f), \tag{20}$$

because their numerators are, respectively, $\Delta\gamma(1-\gamma_f-\Delta\gamma)$, $\Delta\gamma(1-\gamma_f)$, $\Delta\gamma(1-\gamma_f+\Delta\gamma)$, and $\Delta\gamma(1-\gamma_f+2\Delta\gamma)$. The frontier $\theta_{PC_m}(\overline{U}_f)$ is positive only under the following assumption. Under this assumption, all four frontiers tend to $+\infty$ as $\overline{U}_f \to \overline{U}_m^k$.

Assumption 1 *We assume that $\gamma_f < \frac{1+\gamma_m}{2}$, implying that $\gamma_f \in [\gamma_m, \frac{1+\gamma_m}{2})$ and $\Delta\gamma < 1-\gamma_f$.*

Assumption 1 requires heterogeneity in family duties to be sufficiently small. As we show below, it is the condition under which Regime $R1$ exists.

Proposition 3 ***Optimal menu when both genders are employed.*** *Let $\overline{U}_f < \overline{U}_m^k$. The solution to problem (13) is as follows.*

Regime *$R1$, for $\theta \le \theta_{PC_m}(\overline{U}_f)$ (which requires Assumption 1): $IC_m(\theta)$ and $PC_f(\theta)$ bind; men's effort is efficient, $e_m = e_m^B(\theta)$, while women's effort is*

$$e_f^{R1}(\theta) = \frac{1-2\gamma_f+\gamma_m}{\theta} = \frac{1-\gamma_f}{\theta} - \frac{\Delta\gamma}{\theta} < e_f^B(\theta); \tag{21}$$

women receive $U_f = \overline{U}_f$ and men receive $U_m = \overline{U}_f + \Delta\gamma e_f^{R1}(\theta) \ge \overline{U}_m^k$, i.e., the information rent $\Delta\gamma e_f^{R1}(\theta) - \Delta\overline{U}$.

Regime *$R2$, for $\theta_{PC_m}(\overline{U}_f) \le \theta \le \theta^B_{IC_m}(\overline{U}_f)$: $IC_m(\theta)$, $PC_f(\theta)$, and $PC_m(\theta)$ bind; $e_m = e_m^B(\theta)$, while women's effort is*

$$e_f^{R2} = \frac{\overline{U}_m^k - \overline{U}_f}{\gamma_f - \gamma_m}, \qquad \text{with } e_f^{R1}(\theta) \le e_f^{R2} \le e_f^B(\theta); \tag{22}$$

both genders receive their outside options.

Regime *B, for $\theta^B_{IC_m}(\overline{U}_f) \le \theta \le \theta^B_{IC_f}(\overline{U}_f)$: only the participation constraints bind; the benchmark contracts are offered.*

Regime *$R4$, for $\theta^B_{IC_f}(\overline{U}_f) \le \theta \le \theta_{PC_f}(\overline{U}_f)$: $IC_f(\theta)$, $PC_f(\theta)$, and $PC_m(\theta)$ bind; $e_f = e_f^B(\theta)$, while men's effort is*

$$e_m^{R4} = \frac{\overline{U}_m^k - \overline{U}_f}{\gamma_f - \gamma_m}, \qquad \text{with } e_m^B(\theta) \le e_m^{R4} \le e_m^{R5}(\theta); \tag{23}$$

both genders receive their outside options.

Regime *$R5$, for $\theta \ge \theta_{PC_f}(\overline{U}_f)$: $IC_f(\theta)$ and $PC_m(\theta)$ bind; $e_f = e_f^B(\theta)$, while men's effort is*

$$e_m^{R5}(\theta) = \frac{1-2\gamma_m+\gamma_f}{\theta} = \frac{1-\gamma_m}{\theta} + \frac{\Delta\gamma}{\theta} > e_m^B(\theta); \tag{24}$$

Table 1: Summary of regimes under gender-neutral contracting

Regime	Binding constraints	Mimicker	Effort	Rent	Additional exclusion
R1	PC_f, IC_m	Men	$e_f \downarrow$ $e_m = e_m^B$	Men	Women
R2	PC_f, PC_m, IC_m	Men	$e_f \downarrow$ $e_m = e_m^B$	None	Women
B	PC_f, PC_m	**None**	**No distortions**	**None**	**None**
R4	PC_f, PC_m, IC_f	Women	$e_m \uparrow$ $e_f = e_f^B$	None	Men
R5	PC_m, IC_f	Women	$e_m \uparrow$ $e_f = e_f^B$	Women	Men

Notes. In regime B, the gender-neutral menu implements the benchmark allocation. The arrows indicate effort distortions relative to the tagging benchmark. The term "mimicker" denotes the gender whose incentive constraint binds when both genders are employed. "Additional exclusion" identifies the gender whose employment cutoff, $\widetilde{\theta}_g(\overline{U}_f)$, is downward distorted relative to tagging when the marginal worker belongs to the relevant regime. Some workers who would be employed under tagging are therefore excluded under gender neutrality (Proposition 5).

men receive $U_m = \overline{U}_m^k$ *and women receive* $U_f = \overline{U}_m^k - \Delta\gamma e_m^{R5}(\theta) \geq \overline{U}_f$*, i.e., the information rent* $\Delta\overline{U} - \Delta\gamma e_m^{R5}(\theta)$*.*

In all regimes, effort levels and utilities are continuous in θ *and* $\overline{U}_f$*, wages are given by* $w_g = U_g + \frac{1}{2}\theta e_g^2 + \gamma_g e_g$*, and the resulting profit is*

$$V(\theta) = \pi_f^B(\theta) + \pi_m^B(\theta) - \underbrace{\sum_{g\in\{f,m\}} \frac{1}{2}\theta\left(e_g - e_g^B(\theta)\right)^2}_{\textit{efficiency loss}} - \underbrace{\sum_{g\in\{f,m\}} \left(U_g - \overline{U}_g\right)}_{\textit{information rents}}. \tag{25}$$

The proof, in Appendix D, verifies the Kuhn–Tucker conditions of problem (13) regime by regime. Because this problem maximizes a concave function subject to linear constraints, these conditions guarantee that each menu is a global optimum. Equation (25) states that, relative to tagging, gender neutrality can reduce profit in only two ways: through an *efficiency loss*, due to distorted effort, and through *information rents*, left to the mimicking gender.

Table 1 summarizes the main insights from the five regimes. Figure 2 provides their graphical illustration in the space $[\overline{U}_f, \theta]$ by completing Figure 1. The economic logic of the regimes is as follows.

Regimes $R2$ *and* $R1$*: men as the mimickers (*IC_m *binding).* At low values of θ, men would take the contract designed for women. The firm's optimal response is to shorten women's hours. In $R2$, it sets female effort at e_f^{R2}, the longest female contract that men do not want when they receive only their outside option: a man taking this contract would gain exactly $\Delta\gamma e_f^{R2} = \Delta\overline{U}$, which just compensates him for giving up his better outside option. No rent is paid, but women's

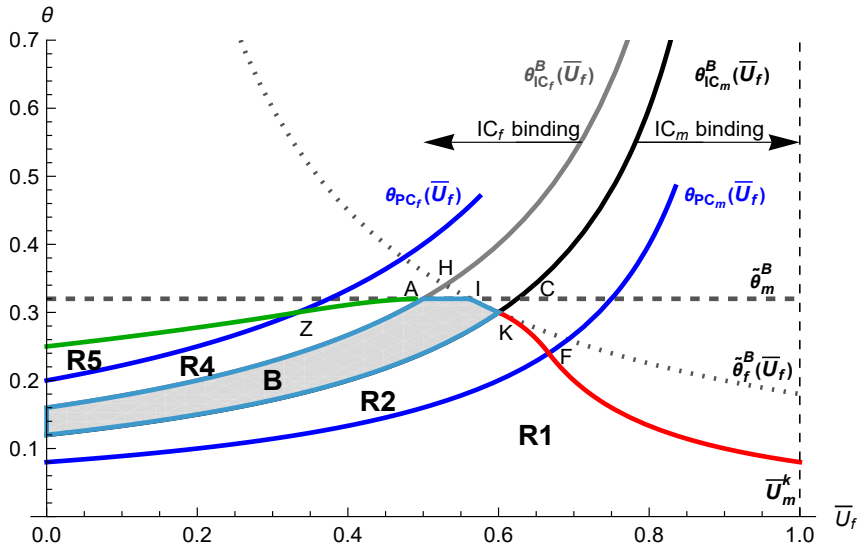


Figure 2: **The five regimes and the employment cutoffs under gender neutrality.** The figure adds to Figure 1 the frontiers $\theta_{PC_f}(\overline{U}_f)$ and $\theta_{PC_m}(\overline{U}_f)$ and the employment cutoffs under gender neutrality. The red curve is the female cutoff $\widetilde{\theta}_f(\overline{U}_f)$ in Regimes $R2$ and $R1$; the green curve is the male cutoff $\widetilde{\theta}_m(\overline{U}_f)$ in Regimes $R5$ and $R4$. Below these curves, both genders are employed. Between the red curve and the dashed line $\widetilde{\theta}^B_m$, only men are employed; between the green curve and the dotted curve $\widetilde{\theta}^B_f(\overline{U}_f)$, only women are employed, in both cases with efficient contracts. Women between the red curve and $\widetilde{\theta}^B_f(\overline{U}_f)$, and men between the green curve and $\widetilde{\theta}^B_m$, are employed under tagging but excluded under gender neutrality. As in Figure 1, $\overline{U}^k_m = 1$, $\gamma_f = 0.4$, and $\gamma_m = 0.2$.

effort is inefficiently low. As θ falls or $\overline{U}_f$ rises, this cap on women's hours becomes more severe, and at $\theta_{PC_m}(\overline{U}_f)$ it becomes cheaper to let men earn a rent than to shorten women's hours further. This is Regime $R1$. There, each additional unit of female effort raises the rent that men obtain by $\Delta\gamma$. The firm therefore chooses female effort as if women's family-duty cost were $\gamma_f + \Delta\gamma$ instead of γ_f, which yields (21). Assumption 1 guarantees that this effort level is positive. Without it, Regime $R1$ does not exist and Regime $R2$ extends to all $\theta \leq \theta^B_{IC_m}(\overline{U}_f)$. Note that $R1$ is the "classical" regime which arises when the incentive constraint of the low-cost type (i.e., male workers) and the participation constraint of the high-cost type are binding. This case has received considerable attention in the contract-theory literature, where outside options are often assumed to be identical across types.

Regimes $R4$ and $R5$: women as the mimickers (IC_f binding). At high values of θ, women would take the contract designed for men. The firm now makes the male contract less attractive to women by lengthening men's hours, since each additional hour is more costly for women than for men (countervailing incentives). In $R4$, it sets male effort at e^{R4}_m, the shortest male contract that women do not want when they receive only their outside option. In $R5$, it becomes cheaper to leave women a rent, and the firm chooses male effort as if men's family-duty cost were $\gamma_m - \Delta\gamma$, which yields (24).

Finally, effort levels move continuously across regimes. At fixed $\overline{U}_f$, female effort equals

$e_f^B(\theta)$ at $\theta_{IC_m}^B(\overline{U}_f)$ and $e_f^{R1}(\theta)$ at $\theta_{PC_m}(\overline{U}_f)$, because e_f^{R2} coincides with these two values at the respective frontiers. The same holds for male effort at $\theta_{IC_f}^B(\overline{U}_f)$ and $\theta_{PC_f}(\overline{U}_f)$.

4.4 Step 3: Who is employed? The extensive margin

Net contributions under gender neutrality and the pointwise hiring rule. The extensive margin is governed by a simple comparison. When one gender's contract is mimicked, employing that gender creates its benchmark profit but also imposes an incentive cost on the contract offered to the other gender. We call the resulting profit gain its *net contribution* under gender neutrality.

In Regimes $R1$ and $R2$, men mimic the female contract, men's effort is efficient, and women receive their outside option. By (25), the value of employing both genders is $V(\theta) = \pi_m^B(\theta) + \pi_f^{GN}(\theta)$, where

$$\pi_f^{GN}(\theta) \equiv \pi_f^B(\theta) - \frac{1}{2}\theta\left(e_f - e_f^B(\theta)\right)^2 - \left(U_m - \overline{U}_m^k\right) \tag{26}$$

is the net contribution of women to the firm's profit under gender neutrality. It equals the profit that women would generate under tagging minus the cost that their employment imposes on the firm through men's incentive constraint: the distortion of women's effort and, in $R1$, the rent paid to men. Substituting women's effort and men's utility in Regimes $R1$ and $R2$ from Proposition 3 into (26) gives

$$\pi_f^{GN}(\theta) = \begin{cases} \dfrac{(1-2\gamma_f+\gamma_m)^2}{2\theta} - \left(2\overline{U}_f - \overline{U}_m^k\right) & \text{if } \theta \le \theta_{PC_m}(\overline{U}_f) \ \ (R1), \\ (1-\gamma_f)\dfrac{\Delta\overline{U}}{\Delta\gamma} - \dfrac{1}{2}\theta\left(\dfrac{\Delta\overline{U}}{\Delta\gamma}\right)^2 - \overline{U}_f & \text{if } \theta_{PC_m}(\overline{U}_f) \le \theta \le \theta_{IC_m}^B(\overline{U}_f) \ \ (R2). \end{cases} \tag{27}$$

Symmetrically, in Regimes $R4$ and $R5$, women's effort is efficient and men receive their outside option, so that $V(\theta) = \pi_f^B(\theta) + \pi_m^{GN}(\theta)$, where

$$\pi_m^{GN}(\theta) \equiv \pi_m^B(\theta) - \frac{1}{2}\theta\left(e_m - e_m^B(\theta)\right)^2 - \left(U_f - \overline{U}_f\right) \tag{28}$$

is the *net contribution of men* under gender neutrality. Substituting men's effort and women's utility in Regimes $R4$ and $R5$ from Proposition 3 into (28) gives

$$\pi_m^{GN}(\theta) = \begin{cases} (1-\gamma_m)\dfrac{\Delta\overline{U}}{\Delta\gamma} - \dfrac{1}{2}\theta\left(\dfrac{\Delta\overline{U}}{\Delta\gamma}\right)^2 - \overline{U}_m^k & \text{if } \theta_{IC_f}^B(\overline{U}_f) \le \theta \le \theta_{PC_f}(\overline{U}_f) \ \ (R4), \\ \dfrac{(1-2\gamma_m+\gamma_f)^2}{2\theta} - \left(2\overline{U}_m^k - \overline{U}_f\right) & \text{if } \theta \ge \theta_{PC_f}(\overline{U}_f) \ \ (R5). \end{cases} \tag{29}$$

The following lemma establishes the properties of the net contributions that lead directly to the firm's hiring rule. First, each net contribution is continuous and strictly decreasing in θ, which generates the cutoff structure of employment. Second, it never exceeds the corresponding benchmark profit, because employing a gender whose contract is mimicked creates additional

incentive costs. Finally, a positive net contribution from that gender implies that employing the other gender is also profitable. These properties allow the firm's hiring decision to be characterized simply by the sign of the relevant employment contribution. The proof is in Appendix E.

Lemma 3 ***Net contributions.*** *(i) In Regimes R1 and R2, $\pi_f^{GN}(\theta)$ is continuous and strictly decreasing in θ. Moreover, $\pi_f^{GN}(\theta) \leq \pi_f^B(\theta)$, with equality at $\theta = \theta^B_{IC_m}(\overline{U}_f)$, and $\pi_f^{GN}(\theta) > 0$ implies $\pi_m^B(\theta) > 0$.*
(ii) In Regimes R4 and R5, $\pi_m^{GN}(\theta)$ is continuous and strictly decreasing in θ. Moreover, $\pi_m^{GN}(\theta) \leq \pi_m^B(\theta)$, with equality at $\theta = \theta^B_{IC_f}(\overline{U}_f)$, and $\pi_m^{GN}(\theta) > 0$ implies $\pi_f^B(\theta) > 0$.

To see the hiring rule at a glance, define the effective contribution of each gender by

$$\pi_f^{\text{emp}}(\theta) \equiv \begin{cases} \pi_f^{GN}(\theta), & \theta \leq \theta^B_{IC_m}(\overline{U}_f), \\ \pi_f^B(\theta), & \theta \geq \theta^B_{IC_m}(\overline{U}_f), \end{cases} \qquad \pi_m^{\text{emp}}(\theta) \equiv \begin{cases} \pi_m^B(\theta), & \theta \leq \theta^B_{IC_f}(\overline{U}_f), \\ \pi_m^{GN}(\theta), & \theta \geq \theta^B_{IC_f}(\overline{U}_f). \end{cases} \tag{30}$$

Thus, a gender's benchmark profit is relevant when its contract is not mimicked, and its effort remains efficient. When its contract is mimicked, the relevant object is instead its net contribution, which deducts from benchmark profit the cost of the effort distortion and any rent paid to the mimicking gender. The two pieces coincide at each frontier by Lemma 3.

At a zero contribution, we break the firm's indifference in favor of fewer employed workers. An unassigned worker who is indifferent between a contract and the outside option selects the outside option. These conventions affect only a measure-zero set of types.

Proposition 4 ***Employment under gender neutrality.*** *Let $\overline{U}_f < \overline{U}_m^k$. For every θ and each gender $g \in \{f, m\}$, the firm employs gender g if and only if*

$$\pi_g^{\text{emp}}(\theta) > 0. \tag{31}$$

Accordingly:
(i) if both employment contributions are positive, the firm employs both genders and offers the menu characterized in Proposition 3;
(ii) if only one employment contribution is positive, the firm employs only that gender, offers its efficient contract, and leaves it exactly its outside-option utility;
(iii) if neither employment contribution is positive, nobody is employed.
In all cases, a gender that is not employed selects its outside option.

The proof in Appendix E uses one argument for both genders. Let g denote the gender whose contract is mimicked and g' the other gender. The profit from employing both is $\pi_{g'}^B(\theta) + \pi_g^{GN}(\theta)$. Starting from employment of g' alone, the gain from adding g is exactly $\pi_g^{GN}(\theta)$. Hence the firm employs g if and only if this net contribution is positive. If it is not, excluding g removes the incentive problem and leaves g' with its efficient contract.

For example, in $R1$ and $R2$, the female contract is mimicked. The firm therefore employs women only when their benchmark profit covers both the effort distortion and, in $R1$, the rent paid to men. The same logic applies to men in $R4$ and $R5$. Thus, additional exclusion always falls on the gender whose contract is mimicked. In the Benchmark region, no contract is mimicked and (31) reduces to the tagging rule. In Figure 1, this gives the female-only region AIH and the male-only region ICK.

Employment cutoffs. Equation (30) defines the employment contribution $\pi_g^{\text{emp}}(\theta)$ by selecting, in each region, the profit that is relevant for the firm's hiring decision. For a given $\overline{U}_f$, its expression changes at the relevant incentive-compatibility frontier. At that frontier, the net contribution $\pi_g^{GN}(\theta)$ coincides with the benchmark profit $\pi_g^B(\theta)$, so the piecewise-defined employment contribution has no jump. Since its two pieces are strictly decreasing in θ and coincide where they meet, $\pi_g^{\text{emp}}(\theta)$ is continuous and strictly decreasing over its entire domain.

Moreover, the employment contribution is positive for sufficiently small θ and negative for sufficiently large θ; Appendix F.2 establishes these limiting properties explicitly. It therefore crosses zero exactly once. Let $\widetilde{\theta}_g(\overline{U}_f)$ denote this unique zero. Then

$$\pi_g^{\text{emp}}(\theta) > 0 \quad \Longleftrightarrow \quad \theta < \widetilde{\theta}_g(\overline{U}_f).$$

By Proposition 4, workers of gender g are therefore employed if and only if their effort cost lies below $\widetilde{\theta}_g(\overline{U}_f)$. The intuition is straightforward: as θ increases, effort becomes more costly and the contribution generated by the worker declines. Once the employment contribution reaches zero, all workers of that gender with higher effort costs are excluded.

The form of the zero depends on the regime containing the marginal worker. Points K and A, introduced in Figure 1, determine whether the female and male cutoffs, respectively, differ from their tagging levels. Two additional points determine which distorted regime contains the zero. Point F is where the female cutoff crosses $\theta_{PC_m}(\overline{U}_f)$, and point Z is where the male cutoff crosses $\theta_{PC_f}(\overline{U}_f)$. Their horizontal coordinates are

$$\overline{U}_f^F = \overline{U}_m^k \frac{1-\gamma_m}{1+2\gamma_f-3\gamma_m}, \qquad \overline{U}_f^Z = \overline{U}_m^k \frac{1-3\gamma_f+2\gamma_m}{1-\gamma_f}. \tag{32}$$

Appendix F.1 derives both points and their vertical coordinates; Appendix C already derived A, I, and K. The relevant ordering is

$$\overline{U}_f^Z < \overline{U}_f^A < \overline{U}_f^I < \overline{U}_f^K < \overline{U}_f^F < \overline{U}_m^k.$$

Proposition 5 ***Marginal workers under gender neutrality.*** *Under Assumption 1:*

(i) The female employment cutoff is

$$\widetilde{\theta}_f(\overline{U}_f) = \begin{cases} \widetilde{\theta}_f^B(\overline{U}_f) = \dfrac{1}{2}\dfrac{(1-\gamma_f)^2}{\overline{U}_f} & \text{if } \overline{U}_f \le \overline{U}_f^K, \\ \widetilde{\theta}_f^{R2}\left(\overline{U}_f\right) = 2\dfrac{\gamma_f-\gamma_m}{\left(\overline{U}_m^k-\overline{U}_f\right)^2}\left(\overline{U}_m^k(1-\gamma_f)-\overline{U}_f(1-\gamma_m)\right) & \text{if } \overline{U}_f^K \le \overline{U}_f \le \overline{U}_f^F, \\ \widetilde{\theta}_f^{R1}\left(\overline{U}_f\right) = \dfrac{1}{2}\dfrac{(1-2\gamma_f+\gamma_m)^2}{2\overline{U}_f-\overline{U}_m^k} & \text{if } \overline{U}_f^F \le \overline{U}_f \le \overline{U}_m^k. \end{cases} \tag{33}$$

(ii) The male employment cutoff is

$$\widetilde{\theta}_m(\overline{U}_f) = \begin{cases} \widetilde{\theta}_m^{R5}\left(\overline{U}_f\right) = \dfrac{1}{2}\dfrac{(1-2\gamma_m+\gamma_f)^2}{2\overline{U}_m^k-\overline{U}_f} & \text{if } \overline{U}_f \le \overline{U}_f^Z, \\ \widetilde{\theta}_m^{R4}\left(\overline{U}_f\right) = 2\dfrac{\gamma_f-\gamma_m}{\left(\overline{U}_m^k-\overline{U}_f\right)^2}\left(\overline{U}_m^k(1-\gamma_f)-\overline{U}_f(1-\gamma_m)\right) & \text{if } \overline{U}_f^Z \le \overline{U}_f \le \overline{U}_f^A, \\ \widetilde{\theta}_m^B = \dfrac{1}{2}\dfrac{(1-\gamma_m)^2}{\overline{U}_m^k} & \text{if } \overline{U}_f \ge \overline{U}_f^A. \end{cases} \tag{34}$$

Both cutoffs are continuous in $\overline{U}_f$. For each value of $\overline{U}_f$, workers of gender g are employed if and only if $\theta < \widetilde{\theta}_g(\overline{U}_f)$.

(iii) Additional exclusion. *Gender neutrality excludes women who would be employed under tagging if and only if $\overline{U}_f > \overline{U}_f^K$: then $\widetilde{\theta}_f(\overline{U}_f) < \widetilde{\theta}_f^B(\overline{U}_f)$. It excludes men who would be employed under tagging if and only if $\overline{U}_f < \overline{U}_f^A$: then $\widetilde{\theta}_m(\overline{U}_f) < \widetilde{\theta}_m^B$.*

(iv) Gender composition. *As under tagging, when both cutoffs lie in the interior of the support $[\underline{\theta}, \overline{\theta}]$, the firm employs more men than women if and only if $\overline{U}_f > \overline{U}_f^I$.*

The proof is in Appendix F. The formulas follow directly from the unique-zero argument. Outside the distorted region, the zero is the tagging cutoff. Inside it, the zero solves $\pi_f^{GN}(\theta) = 0$ for women or $\pi_m^{GN}(\theta) = 0$ for men. Points K and A mark the switch between these two cases; points F and Z mark the switch between the two formulas for the relevant net contribution. For example, between K and F, the marginal woman lies in $R2$, so $\tilde{\theta}_f^{R2}(\overline{U}_f)$ is obtained by setting the second line of (27) equal to zero.

The piece $\widetilde{\theta}_m^{R5}(\overline{U}_f)$ is relevant only if $\overline{U}_f^Z > \overline{U}_f^o > 0$. In particular, $\overline{U}_f^Z$ must be positive, which, by (32), requires $\gamma_f < (1+2\gamma_m)/3$. Hence, Regime $R5$ with both genders employed can exist only under the following assumption.

Assumption 2 *We assume that $\gamma_f < \frac{1+2\gamma_m}{3} < \frac{1+\gamma_m}{2}$. Consequently, $\gamma_f \in [\gamma_m, (1+2\gamma_m)/3)$ and $\Delta\gamma < (1-\gamma_m)/3$.*

Assumption 2 is more stringent than Assumption 1. When it fails, men's net contribution is negative whenever the firm would have to leave a rent to women, so men are excluded before Regime $R5$ is reached.

Figure 2 illustrates the employment cutoffs generated by the binding incentive constraints. The red curve is the female cutoff for $\overline{U}_f \geq \overline{U}_f^K$. It starts at point K, where the gender-neutral and benchmark cutoffs coincide, follows $\widetilde{\theta}_f^{R2}(\overline{U}_f)$ until it meets $\theta_{PC_m}(\overline{U}_f)$ at point F, and then follows $\widetilde{\theta}_f^{R1}(\overline{U}_f)$. Symmetrically, the green curve is the male cutoff for $\overline{U}_f \leq \overline{U}_f^A$. It follows $\widetilde{\theta}_m^{R5}(\overline{U}_f)$ up to point Z, and then $\widetilde{\theta}_m^{R4}(\overline{U}_f)$ until it reaches the benchmark cutoff $\widetilde{\theta}_m^B$ at point A. The vertical distance between the red curve and $\widetilde{\theta}_f^B(\overline{U}_f)$, and between the green curve and $\widetilde{\theta}_m^B$, identifies workers who are profitable under tagging but excluded under gender neutrality. The figure illustrates a general principle: the distorted cutoff belongs to the gender whose contract is being mimicked. This principle is summarized in the final column of Table 1.

Comparative statics of additional exclusion. The response to a higher female outside option can already be seen from the net contributions:

$$\frac{\partial \pi_f^{GN}(\theta)}{\partial \overline{U}_f} = -(1+\mu_m) < 0, \qquad \frac{\partial \pi_m^{GN}(\theta)}{\partial \overline{U}_f} = \mu_f \geq 0, \tag{35}$$

where μ_m and μ_f are the non-negative Lagrange multipliers of $IC_m(\theta)$ and $IC_f(\theta)$ in problem (13); see Appendix D. Thus, a higher $\overline{U}_f$ lowers the net contribution of women when men mimic their contract, but raises the net contribution of men when women mimic theirs.

To compare exclusion across genders, define the proportional extensive-margin distortion as $1 - \tilde{\theta}_g(\overline{U}_f)/\tilde{\theta}_g^B$. This measure is zero when gender neutrality leaves the cutoff unchanged and is positive when it induces additional exclusion.

Corollary 1 ***Comparative statics of additional exclusion.*** *As $\overline{U}_f$ increases:*

(i) for $\overline{U}_f \geq \overline{U}_f^K$, the female cutoff and the ratio $\tilde{\theta}_f(\overline{U}_f)/\tilde{\theta}_f^B(\overline{U}_f)$ are strictly decreasing. The ratio equals one at K, so the proportional exclusion of women increases beyond K and is larger when the marginal woman lies in $R1$ than when she lies in $R2$;

(ii) for $\overline{U}_f \leq \overline{U}_f^A$, the male cutoff and the ratio $\tilde{\theta}_m(\overline{U}_f)/\tilde{\theta}_m^B$ are strictly increasing. The ratio reaches one at A, so the proportional exclusion of men decreases up to A and is larger when the marginal man lies in $R5$ than when he lies in $R4$.

The proof is in Appendix F. The two parts have the same economic logic. A higher $\overline{U}_f$ tightens men's incentive constraint when the female contract is mimicked, so women become more costly to add and their cutoff moves farther below its tagging level. When the male contract is mimicked, the same increase makes that contract less attractive to women; the incentive cost falls and the male cutoff returns toward its tagging level.[9]

The additional exclusion of men in $R5$ is not contradicted by their upward effort distortion. Moving effort either above or below its efficient level is costly. Excluding men eliminates both

[9]For women, the absolute gap $\tilde{\theta}_f^B - \tilde{\theta}_f$ need not be monotone in $R1$ because the tagging cutoff itself changes with $\overline{U}_f$. The ratio in the corollary isolates the proportional distortion created by gender neutrality.

this cost and the rent paid to women. Figure 2 therefore has a symmetric interpretation: the green distortion is largest at low $\overline{U}_f$ and vanishes at A, whereas the red distortion begins at K and increases as the two outside options converge.

4.5 Existence of the different Regimes

The size and existence of the different regimes depend on the workers' heterogeneity, that is, the difference in the additional cost of providing labor supply associated with family duties, $\Delta\gamma$. As $\Delta\gamma$ becomes larger, the size of the Benchmark region, as well as the regions where Regimes $R2$ and $R4$ occur, expands, whereas Regimes $R1$ and $R5$ may cease to exist.

Formally, from the results stated so far, we can derive the following corollary.

Corollary 2 ***Existence and shape of the different regimes.*** *(i) The horizontal gaps among the curves $\theta_{PC_f}(\overline{U}_f)$, $\theta^B_{IC_f}(\overline{U}_f)$, $\theta^B_{IC_m}(\overline{U}_f)$, and $\theta_{PC_m}(\overline{U}_f)$ widen as $\Delta\gamma$ increases. In particular, the region in which benchmark contracts are incentive compatible expands. (ii) Regime $R1$ exists if and only if Assumption 1 holds; Regime $R5$ with both genders employed can exist only under the tighter bound stated in Assumption 2. (iii) As the outside option of female workers increases, the sequence of regimes is $R5$ (if it exists), $R4$, the Benchmark, $R2$, and $R1$ (if it exists). (iv) For any given $\overline{U}_f$, the regime depends on θ: as θ increases, the menu offered when both genders are employed moves from $R1$ (if it exists) to $R2$, the Benchmark, $R4$, and $R5$. Whether both genders are actually employed depends on the cutoffs of Proposition 5. If $\overline{U}_f$ is sufficiently small, the whole sequence can be observed with both genders employed. At higher values of $\overline{U}_f$, some regimes disappear because workers with higher values of θ are excluded before those regimes are reached. Eventually, regimes in which both genders are employed may cease to exist.*

Let us start from Point (i). At a given θ, the horizontal distance between two adjacent frontiers is $(\Delta\gamma)^2/\theta$ (see (19) and (18)). Intuitively, as workers' heterogeneity in the additional cost of family duties increases, mimicking becomes less appealing for the workers, making incentive compatibility more likely. Point (ii) holds because, when heterogeneity is large, leaving a rent to the mimicking gender becomes too costly relative to the alternatives. In $R1$, the effort that the firm would assign to women, $e_f^{R1}(\theta)$, is no longer positive when Assumption 1 fails, so that the firm always prefers to cap women's hours as in $R2$. In $R5$, when Assumption 2 fails, men's net contribution is negative whenever the firm would have to leave a rent to women, so that men are excluded before $R5$ is reached. The remaining points are most easily understood from Figure 2.[10] To illustrate Point (iv) of Corollary 2, consider the intermediate value $\overline{U}_f = 0.5$ in Figure 2. Holding the female workers' outside option fixed at this value, an increase in θ moves the relevant regime from $R1$ to $R2$ and then to the Benchmark, B. More generally, gender neutrality

[10] Figure 2 illustrates a case in which both $R1$ and $R5$ exist. As $\Delta\gamma$ increases, the horizontal gaps among the vertical curves $\theta_{PC_f}(\overline{U}_f)$, $\theta^B_{IC_f}(\overline{U}_f)$, $\theta^B_{IC_m}(\overline{U}_f)$, and $\theta_{PC_m}(\overline{U}_f)$ widen, and regimes $R1$ and $R5$ eventually cease to exist.

affects workers differently depending on their value of θ. Because Regime $R1$ applies to the most efficient workers, those with low effort costs (see Figure 2), it is the effort and labor income of these women that are most affected by gender neutrality, although their utility remains equal to their outside option.

Finally, Regime $R5$ is the only regime in which female workers receive a rent and, therefore, the only one in which gender neutrality (GN) benefits them. However, this regime requires both a sufficiently small gender gap in γ (Assumption 2) and a sufficiently large gender gap in outside options $\overline{U}$. These conditions are unlikely to hold simultaneously. Indeed, the same cultural and historical forces that limit women's labor-market opportunities also tend to increase their family responsibilities, and hence γ_f. Thus, the two dimensions of heterogeneity, $\Delta\gamma$ and $\Delta\overline{U}$, are likely to be positively correlated.[11] Consequently, Regime $R5$, where GN raises women's utility, is less likely to occur than $R1$, where GN leaves women's utility unchanged at their outside option $\overline{U}_f$ but reduces female employment, effort, and labor income.

4.6 Effort, wages, and workers' utilities across regimes

We illustrate the implications of the different regimes by fixing $\theta = 0.21$, a value for which all five regimes arise in the example represented in Figure 2. As the female outside option $\overline{U}_f$ increases, the economy moves from Regime R5 to R4, then to the Benchmark region B, and finally to R2 and R1. This exercise allows us to compare the gender-neutral allocation with the tagging benchmark. Under tagging, effort levels are efficient and independent of $\overline{U}_f$; as a result, the gender wage gap decreases monotonically as women's outside options improve, but women always provide lower effort than men and thus earn less. Under gender neutrality, by contrast, incentive compatibility constraints introduce effort distortions and, in some regimes, information rents. Consequently, changes in $\overline{U}_f$ affect wages not only through reservation utilities, but also through the screening distortions needed to sustain a gender-neutral menu.

4.6.1 Effort distortions

Figure 3 reports male and female effort levels across the five regimes. In Regime $R1$, women cease to be employed under gender neutrality when the female outside option reaches approximately 0.69. Under tagging, by contrast, the same women would remain employed until $\overline{U}_f$ reaches approximately 0.86. The interval between these two thresholds, therefore, represents additional female exclusion generated by the gender-neutrality constraint.

In the tagging benchmark, effort levels are efficient and therefore constant with respect to $\overline{U}_f$: male workers supply more effort than female workers because they face a lower family-duty cost and obtain higher labor income. Under gender neutrality, this pattern is preserved only in

[11] Furthermore, recall that Assumption 2, which is necessary for the existence of Regime R5 with both genders employed, is more stringent than Assumption 1, which is required for Regime R1.

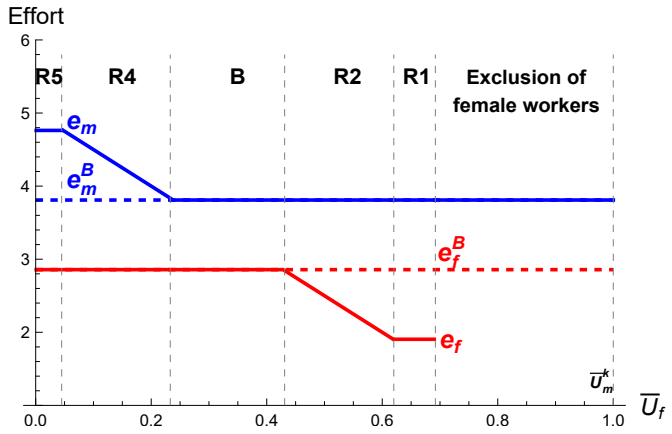


Figure 3: **Effort levels (labor supplies).** Male and female effort levels, e_f and e_m, in the different regimes when $\theta = 0.21$. Benchmark efforts are the dotted horizontal lines e_f^B and e_m^B. As in the previous figures, $\overline{U}_m^k = 1$, $\gamma_f = 0.4$, and $\gamma_m = 0.2$.

the Benchmark region B, where both incentive constraints are slack and the efficient allocation remains implementable.

Outside the Benchmark region, the binding incentive constraint determines which effort level is distorted. In Regimes R5 and R4, female workers are the mimickers, so the firm distorts male effort upward in order to make the male contract less attractive to women. Female effort remains efficient. Conversely, in Regimes R2 and R1, male workers are the mimickers, so the firm distorts female effort downward in order to reduce the attractiveness of the female contract to men. Male effort remains efficient. Hence, gender neutrality affects labor supply asymmetrically across regimes: it increases male effort when IC_f binds, and reduces female effort when IC_m binds. In all regimes, effort distortions translate into larger effort gaps.

4.6.2 Wage-gap implications

Figure 4 shows the gender wage gap under gender neutrality. In the tagging benchmark, effort levels are efficient, and both participation constraints bind, so that an increase in $\overline{U}_f$ directly narrows the wage gap, but women always provide lower labor supply and earn less. Under gender neutrality, instead, wages reflect not only outside options, but also the effort distortions required by incentive compatibility and, in some regimes, information rents. Hence, the gender wage gap is no longer monotonically decreasing in the female outside option $\overline{U}_f$.

In the central regimes, the gender wage gap is given by

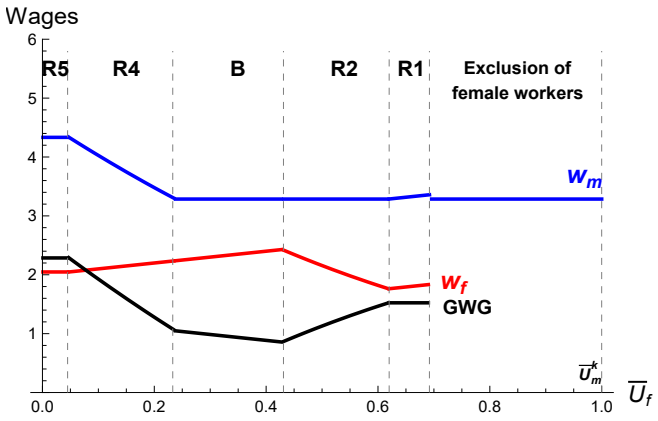


Figure 4: **The Gender Wage Gap** (GWG). Male and female wages and the gender wage gap in the different regimes when $\theta = 0.21$. As in the previous figures, $\overline{U}_m^k = 1$, $\gamma_f = 0.4$, and $\gamma_m = 0.2$.

$$GWG^{R4} = \frac{\left[(1+\gamma_f)(\gamma_f-\gamma_m)+\theta(\overline{U}_m^k-\overline{U}_f)\right]\left[(1-\gamma_f)(\gamma_m-\gamma_f)+\theta(\overline{U}_m^k-\overline{U}_f)\right]}{2\theta(\gamma_f-\gamma_m)^2},$$

$$GWG^{B} = \overline{U}_m^k - \overline{U}_f + \frac{(\gamma_f-\gamma_m)(\gamma_f+\gamma_m)}{2\theta},$$

$$GWG^{R2} = \frac{(\gamma_f-\gamma_m)^2 - \left[(\gamma_f-\gamma_m)\gamma_m+\theta(\overline{U}_m^k-\overline{U}_f)\right]^2}{2\theta(\gamma_f-\gamma_m)^2}.$$

In Regimes $R1$ and $R5$, instead, the gender wage gap is independent of $\overline{U}_f$:

$$GWG^{R1} = \frac{2(\gamma_f-\gamma_m)(1-\gamma_f+\gamma_m)}{\theta}, \qquad GWG^{R5} = \frac{2(\gamma_f-\gamma_m)(1+\gamma_f-\gamma_m)}{\theta}.$$

These expressions show that the gender wage gap decreases with $\overline{U}_f$ in Regime $R4$ and in the Benchmark region B, but increases with $\overline{U}_f$ in Regime $R2$. The gender wage gap therefore reaches its minimum at the boundary between B and $R2$: at this point, the outside-option gap is relatively small and effort levels are still undistorted.

Moving further into $R2$, an increase in $\overline{U}_f$ strengthens the downward distortion of female effort needed to deter male mimicking. This lowers the wage required to deliver women their reservation utility and offsets the direct effect of the higher outside option. As a result, the gender wage gap increases with $\overline{U}_f$ in $R2$; see Figure 4. Gender neutrality therefore generates an additional source of wage inequality: even as women's outside options improve, the screening distortions induced by a gender-neutral menu widen the wage gap relative to the tagging benchmark.

Note that w_g denotes total labor income, while hourly pay is given by w_g/e_g. Hourly wages are much more similar across genders: in our example, the hourly wage gap narrows as $\overline{U}_f$ increases and is nearly zero in Regime R2. Thus, under gender neutrality, the widening gender

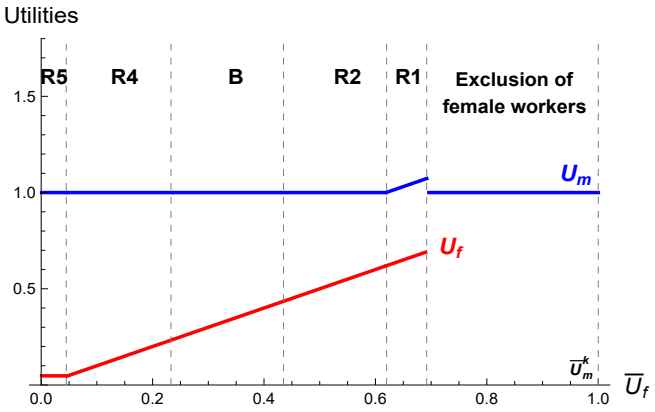


Figure 5: **Utilities.** Utilities of employed workers in the different regimes when $\theta = 0.21$. As in the previous figures, $\overline{U}_m^k = 1$, $\gamma_f = 0.4$, and $\gamma_m = 0.2$.

gap in labor income is driven primarily by differences in hours worked, while differences in hourly pay play only a minor role.

4.6.3 Utilities and exclusion

Figure 5 shows that utilities are less affected by gender neutrality than wages. In Regimes $R2$, B, and $R4$, both participation constraints bind, so male and female workers receive their respective outside-option utilities, $\overline{U}_m^k$ and $\overline{U}_f$. Hence, even when wages differ substantially across genders, these wage differences do not translate one-for-one into utility differences: they partly compensate workers for different effort levels and family-duty costs.

Information rents arise only in the two extreme regimes. In Regime $R1$, the participation constraint of male workers is slack, and men receive an information rent as long as both genders are employed. Once women are excluded, this rent disappears, and men's utility returns to $\overline{U}_m^k$. Symmetrically, in Regime $R5$, the participation constraint of female workers is slack, and women receive an information rent as long as both genders are employed. Such rent increases with $\overline{U}_m^k - \overline{U}_f$, and women's total utility is constant within $R5$.

The extensive margin is crucial for understanding the effects of gender neutrality. In the tagging benchmark, an increase in women's outside option already lowers the female-employment cutoff because women become more costly to recruit. In Regime $R1$, however, the incentive-compatible gender-neutral menu lowers this cutoff further. As $\overline{U}_f$ approaches $\overline{U}_m^k$, the firm must either distort contracts more strongly or leave rents to male workers. It therefore ceases to employ some women who would still be profitably employed under tagging. Thus, although improvements in women's outside options reduce the outside-option gap, they also increase female exclusion under gender-neutral contracting. Note, however, that excluded women obtain their outside-option utility, as in Regime $R1$, and are therefore no worse off than in that regime. In our model, exclusion by the firm does not imply unemployment.

A final remark on the extensive margin for men is warranted at this stage. Our discussion in Section 4.6 is based on a specific value of θ such that no male worker is excluded. However, as mentioned in Section 4.5, regimes featuring the employment of both genders may not exist. For instance, consider the case $\theta = 0.31$ in Figure 2. For this level of the effort cost, male workers are fully excluded in Regime $R5$ and partially excluded in Regime $R4$, whereas female workers are fully excluded in Regimes $R2$ and $R1$.

In our illustrative example, women with $\theta = 0.21$ (or higher effort costs) cease to be employed before their utility within the firm can converge to men's utility. After exclusion, they obtain their outside-option utility, $\overline{U}_f$, which converges to $\overline{U}_m^k$ as the gender gap in outside options disappears. When the effort cost is sufficiently low, as in the case $\theta = 0.05$, women remain employed as their outside option approaches that of men. In that case, however, their labor income and utility remain below men's because men receive an information rent in Regime $R1$.

4.7 The cost of gender neutrality for the firm

We now turn to the firm. Under tagging, the firm implements the efficient allocation and leaves each employed worker at her or his outside option. Holding θ fixed, benchmark profit decreases with $\overline{U}_f$ as long as women are employed, because satisfying their participation constraint becomes increasingly costly.

Figure 6 compares the profit generated by workers with $\theta = 0.21$ under gender neutrality, Π^{GN}, and under tagging, Π^B. Gender neutrality weakly reduces profit: $\Pi^{GN} \leq \Pi^B$. The two coincide in the Benchmark region B, where the tagging contracts are incentive compatible, and again when both policies lead to the same exclusion decision. Whenever gender neutrality changes the allocation, the shaded profit gap has two components (see equation (25)). Distorted effort and the exclusion of workers who would be employed under tagging generate an *efficiency loss*; in the two extreme regimes, the firm also leaves an *information rent* to women in $R5$ and to men in $R1$.

The behavior of the profit gap reflects the identity of the mimicking gender. Starting from low values of $\overline{U}_f$, the gap decreases through $R5$ and $R4$: women's information rent falls in $R5$, while the upward distortion of male effort becomes smaller in $R4$. The gap vanishes in region B. To the right of B, the pattern is reversed. As $\overline{U}_f$ increases, the female contract becomes more attractive to men. The firm responds by progressively reducing female effort in $R2$ and, in $R1$, by also leaving men an increasing information rent. The profit gap therefore widens until gender neutrality induces the exclusion of women with $\theta = 0.21$, at $\overline{U}_f \approx 0.69$ in the figure. At this point, women's net contribution $\pi_f^{GN}(\theta)$ becomes zero, and the firm switches to employing men only (Proposition 4).

This peak can be established directly for any fixed θ. Before female exclusion, the female-side profit loss is $\pi_f^B(\theta) - \pi_f^{GN}(\theta)$. Since $\partial \pi_f^B / \partial \overline{U}_f = -1$ and (35) gives $\partial \pi_f^{GN}(\theta) / \partial \overline{U}_f = -(1 + \mu_m)$, its derivative is $\mu_m \geq 0$. After exclusion, the profit loss is the forgone benchmark profit $\pi_f^B(\theta)$,

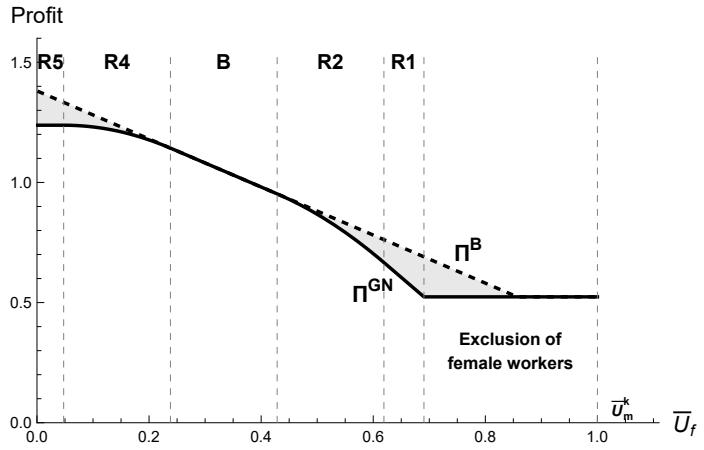


Figure 6: **The firm's profit.** Profit generated by workers with $\theta = 0.21$ under gender neutrality (Π^{GN}, solid line) and under tagging (Π^{B}, dashed line). The shaded area is the firm's profit loss due to gender neutrality. It reflects effort and employment distortions and, in $R1$ and $R5$, the information rent received by the mimicking gender. As in the previous figures, $\overline{U}_m^k = 1$, $\gamma_f = 0.4$, and $\gamma_m = 0.2$.

whose derivative is -1. Hence, for a given type, the loss is maximized at the switch point $\pi_f^{GN}(\theta) = 0$. This is a pointwise result: because different values of θ reach their exclusion thresholds at different values of $\overline{U}_f$, the location of the maximum of the aggregate profit loss would require a separate integration argument.

Beyond this threshold, profit under gender neutrality is the profit from employing men alone. The gap then equals the benchmark profit forgone on women and narrows as $\overline{U}_f$ increases, disappearing at the tagging cutoff, $\overline{U}_f \approx 0.86$, where women of this type would be excluded under either policy. Thus, in this example, the cost of gender neutrality is largest when women's outside options are close to, but still below, men's. This profit loss should not be interpreted as a corresponding loss for women. In $R2$, $R1$, and the subsequent exclusion region, women obtain $\overline{U}_f$ under both tagging and gender neutrality. The information rent in $R1$ is therefore a transfer from the firm to men, whereas the remaining profit gap represents a loss of surplus caused by inefficient effort and the forgone gains from employing women who would still be hired under tagging.

Welfare implications. Taken together, the preceding results distinguish distributional effects from genuine welfare losses. Under tagging, the firm implements the surplus-maximizing allocation and leaves employed workers at their outside-option utilities. Gender neutrality preserves this allocation in region B and wherever both policies implement the same single-gender allocation. In $R2$ and $R4$, workers continue to receive their outside options, so the firm's profit loss corresponds entirely to lower total surplus. In $R1$ and $R5$, part of that loss becomes an information rent for the mimicking gender and is therefore a transfer rather than a welfare loss. The genuine welfare loss arises from distorted effort and from the exclusion of workers whose employment would generate positive surplus under tagging. Consequently, changes in wages

and profits do not translate one-for-one into worker welfare: utility gaps may narrow while employment and labor-income gaps widen, and excluded workers retain their outside-option utility. Nevertheless, total surplus falls whenever gender neutrality generates distortions for a positive mass of workers.

5 Concluding comments

This paper has developed a monopsonistic labor-market model to examine how gender-neutral hiring and compensation policies interact with gendered differences in family-duty costs and outside options. Our analysis shows that, contrary to their formal intention, gender-neutral policies may generate substantial distortions in effort levels, wages, and employment patterns when workers differ in the additional cost of labor associated with family responsibilities. Because firms cannot condition contracts on gender, a common menu distorts effort and employment and may generate information rents for the mimicking gender.

Five distinct regimes may emerge depending on the level of female outside options and the relative size of gendered labor-supply costs. In the central benchmark region—where incentive constraints are slack—the efficient, gender-specific contracts remain implementable despite the gender-neutral restriction. But when incentive constraints bind, effort distortions arise: female effort is downward-distorted when male workers are the mimickers (Regimes $R1$–$R2$), whereas male effort is upward-distorted when women are the mimickers (Regimes $R4$–$R5$). These distortions naturally translate into a non-monotonic gender wage gap and, in Regimes $R1$ and $R5$, in asymmetric information rents which accrue to men when $\overline{U}_m$ is high and to women only in the relatively narrow region where countervailing incentives dominate (Regime $R5$).

A central—and perhaps most striking—implication of the model is that gender-neutral screening can amplify exclusion at the extensive margin. In Regimes $R1$ and $R2$, gender neutrality introduces an additional screening distortion: the female cutoff lies below its tagging benchmark and decreases proportionally faster as women's outside options converge toward those of men. Equal treatment, therefore, generates additional female exclusion relative to the benchmark. Symmetrically, when women's outside options are sufficiently low, gender neutrality lowers the male-employment cutoff below its tagging level and generates additional male exclusion. For a given efficiency type, the firm's profit loss relative to tagging peaks when selective exclusion of that type begins, showing that identical contractual treatment can undermine both efficiency and equality when family responsibilities remain gendered.

A second lesson concerns the role and magnitude of gender heterogeneity in family-duty costs, $\gamma_f - \gamma_m$. When this heterogeneity is large enough for Assumption 2 to fail but still small enough for Assumption 1 to hold, Regime $R5$ cannot arise with both genders employed. Since differences in outside options and differences in family-duty costs are likely positively correlated in practice, our model predicts that the situations in which gender-neutrality helps women are

also precisely those least likely to occur in reality.

Overall, the analysis highlights that *labor-market equality cannot be achieved solely through equal-treatment rules imposed on employers.* When underlying family responsibilities are unequally distributed, gender-neutral screening mechanisms inadvertently reproduce or even amplify gender inequality. This mechanism may also help explain why the positive trend toward narrowing gender gaps in the labor market has slowed down in recent years: as women's outside options improve, the remaining asymmetry in family duties becomes increasingly relevant for firms' screening and hiring decisions. True equal opportunity in the labor market, therefore, requires greater symmetry in household-level constraints: only when family duties are shared more evenly can gender-neutral policies deliver on their promise of fairness.

Introducing employer competition is a natural extension. In a *covered* labor market with two competing firms, competition is likely to mitigate the distortion operating through the extensive margin since a worker not hired by one employer is hired by its rival. In this sense, our monopsonistic setting is particularly useful for highlighting the extensive-margin implications of gender-neutral screening. Under competition, the same underlying force would appear as gendered sorting across firms and contracts. Nevertheless, as long as firms retain some market power and offer self-selecting contracts, gender-neutral screening should still generate gendered outcomes and similar effort distortions because the source of heterogeneity—unequal family-duty costs—remains unchanged. Under perfect competition, profits would be driven to zero, and workers would obtain the entire surplus in equilibrium.

From a policy perspective, our analysis suggests that reducing heterogeneity in family-duty costs is a prerequisite for the effectiveness of gender-neutral labor market regulations. Gender-neutral rules—such as equal pay or pay transparency mandates—are unlikely to deliver equal outcomes unless they are complemented by policies that reduce gender asymmetries in family responsibilities. Moreover, family-friendly arrangements—e.g., remote work—should be designed so as to avoid self-selection along gender lines; otherwise, they may reinforce the correlation between gender and labor-supply costs that underlies inefficient screening.

Policies that increase men's involvement in caregiving play a central role. Beyond compulsory and non-transferable paternity leave, such policies include well-paid individual parental leave entitlements for both parents, and symmetric rights to flexible work arrangements granted by default rather than upon request, which may further contribute to shifting caregiving responsibilities toward a more equal distribution. In the framework of the model, these interventions raise γ_m and reduce the gap $\gamma_f - \gamma_m$. Finally, policies that reduce the private cost of caregiving—such as universal childcare and public long-term care—may directly lower female labor-supply costs, γ_f, thus reducing heterogeneity in family duties.

All these interventions reduce the gender asymmetry in family-duty costs that drives the incentive problem under gender-neutral contracting. In Regime $R1$, this problem affects all relevant margins: it lowers female effort below its tagging level, may require the firm to leave

an information rent to men, and ultimately generates additional female exclusion.

References

Mark Aguiar and Erik Hurst. Measuring trends in leisure: The allocation of time over five decades. *Quarterly Journal of Economics*, 122:969–1006, 2007.

George A. Akerlof. The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning. *American Economic Review*, 68(1):8–19, 1978.

Stefania Albanesi and Claudia Olivetti. Home production, market production and the gender wage gap: Incentives and expectations. *Review of Economic Dynamics*, 12(1):80–107, 2009.

Alberto Alesina, Andrea Ichino, and Loukas Karabarbounis. Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy*, 3(2):1–40, 2011.

Heather Antecol, Kelly Bedard, and Jenna Stearns. Equal but inequitable: Who benefits from gender-neutral tenure clock stopping policies? *American Economic Review*, 108(9):2420–2441, 2018.

Francesca Barigozzi, Cesare Di Timoteo, and Chiara Monfardini. The gender gaps in time-use within italian households during 2002–2014. *Italian Economic Journal*, 9:1263–1296, 2023.

Erling Barth and Harald Dale-Olsen. Monopsonistic discrimination, worker turnover, and the gender wage gap. *Labour Economics*, 16(5):589–597, 2009.

Marianne Bertrand. Gender in the twenty-first century. *AEA Papers and Proceedings*, 110:1–24, 2020.

Suzanne M. Bianchi, Liana C. Sayer, Melissa A. Milkie, and John P. Robinson. Housework: Who did, does or will do it, and how much does it matter? *Social Forces*, 91(1):55–63, 2012.

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

Alessandra Casarico and Salvatore Lattanzio. Behind the child penalty: Understanding what contributes to the labour market costs of motherhood. *Journal of Population Economics*, 36: 1489–1511, 2023.

Helmuth Cremer, Firouz Gahvari, and Jean-Marie Lozachmeur. Tagging and income taxation: Theory and an application. *American Economic Journal: Economic Policy*, 2(1):31–50, 2010.

Daniel Fernández-Kranz and Nuria Rodríguez-Planas. Too family friendly? the consequences of parent part-time working rights. *Journal of Public Economics*, 197:104407, 2021.

Antonio Filippin and Andrea Ichino. Gender wage gap in expectations and realizations. *Labour Economics*, 12(1):125–145, 2005.

Luca Flabbi and Andrea Moro. The effect of job flexibility on female labor market outcomes: Estimates from a search and bargaining model. *Journal of Econometrics*, 168(1):81–95, 2012.

Maria Frech and Gerard Maideu-Morera. The hidden demand for flexibility: A theory for gendered employment dynamics. TSE Working Paper 24–1588, Toulouse School of Economics, 2024.

José Ignacio Giménez-Nadal and Almudena Sevilla. Trends in time allocation: A cross-country analysis. *European Economic Review*, 56(6):1338–1359, 2012.

Claudia Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104 (4):1091–1119, 2014.

Bruno Jullien. Participation constraints in adverse selection models. *Journal of Economic Theory*, 93(1):1–47, 2000.

Henrik Kleven, Camille Landais, Johannes Posch, Andreas Steinhauer, and Josef Zweimüller. Child penalties across countries: Evidence and explanations. *AEA Papers and Proceedings*, 109:122–126, 2019a.

Henrik Kleven, Camille Landais, and Jakob Egholt Søgaard. Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209, 2019b.

Justina Klimaviciute and Pierre Pestieau. The economics of long-term care: An overview. *Journal of Economic Surveys*, 37:1192–1213, 2023.

Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal–Agent Model*. Princeton University Press, Princeton, NJ, 2002.

Tracy R. Lewis and David E. M. Sappington. Countervailing incentives in agency problems. *Journal of Economic Theory*, 49(2):294–313, 1989.

Kjell Erik Lommerud, Odd Rune Straume, and Steinar Vagstad. Mommy tracks and public policy: On self-fulfilling prophecies and gender gaps in hiring and promotion. *Journal of Economic Behavior & Organization*, 116:540–554, 2015.

Alan Manning. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press, Princeton, NJ, 2003.

Alexandre Mas and Amanda Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–3759, 2017.

Barbara Petrongolo. The gender gap in employment and wages. *Nature Human Behaviour*, 3: 316–318, 2019.

Matthew Wiswall and Basit Zafar. Preference for the workplace, investment in human capital, and gender. *Quarterly Journal of Economics*, 133(1):457–507, 2018.

Appendix

A Proof of Lemma 1

Logic of the proof. The only reason why the firm might be unable to obtain the benchmark profit $\pi_g^B(\theta)$ when employing gender g alone is that the benchmark contract may attract workers of the other gender and therefore violate their no-entry constraint. The two points of the lemma deal with this possibility. Point (i) shows that, when the benchmark contract attracts the other gender and generates positive profit, employing both genders dominates employing only gender g. Point (ii) then shows that an optimal single-gender contract never needs to be distorted: either the benchmark contract satisfies the no-entry constraint, or single-gender employment is not optimal.

Point (i). Suppose that $\pi_m^B(\theta) > 0$ and that women strictly prefer the benchmark male contract to their outside option:

$$\overline{U}_m^k - \Delta\gamma e_m^B(\theta) > \overline{U}_f.$$

Offer the single contract $\left(e_m^B(\theta), w_m^B(\theta)\right)$ to both genders. In the notation of problem (13), this allocation gives

$$e_f = e_m = e_m^B(\theta), \qquad U_m = \overline{U}_m^k, \qquad U_f = \overline{U}_m^k - \Delta\gamma e_m^B(\theta).$$

Both participation constraints are satisfied. Moreover, because the two genders choose the same contract, both incentive constraints hold with equality. The pooling allocation is therefore feasible for problem (13), which does not require the contracts chosen by the two genders to differ.

Profit from a given effort–wage pair is independent of gender. Hence each worker generates $\pi_m^B(\theta)$, and the pooling allocation yields $2\pi_m^B(\theta)$. Since $V(\theta)$ is the maximum profit from employing both genders,

$$V(\theta) \geq 2\pi_m^B(\theta) > \pi_m^B(\theta).$$

Moreover, $\pi_m^B(\theta)$ is the unconstrained maximum profit from one male worker. Any contract employing only men therefore generates at most $\pi_m^B(\theta)$. It follows that employing only men cannot be optimal. The argument for women is symmetric.

Point (ii). Let $V_M(\theta)$ and $V_F(\theta)$ denote the maximum profits from employing only men and only women, respectively. Consider a gender $g \in \{f, m\}$ and let $g' \neq g$ denote the other gender. Dropping the no-entry constraint of g' gives the tagging problem of g. Therefore, the benchmark profit is an upper bound on the profit from employing only g:

$$V_g(\theta) \leq \pi_g^B(\theta).$$

There are two cases.

First, suppose that the benchmark contract of gender g satisfies the no-entry constraint of gender g'. The benchmark contract is then feasible when only g is employed. Since it attains the upper bound,

$$V_g(\theta) = \pi_g^B(\theta).$$

Second, suppose that the benchmark contract of gender g violates the no-entry constraint of gender g'. If $\pi_g^B(\theta) > 0$, point (i) gives

$$V(\theta) \geq 2\pi_g^B(\theta) > \pi_g^B(\theta) \geq V_g(\theta).$$

Thus, employing both genders dominates every allocation employing only gender g. If instead $\pi_g^B(\theta) \leq 0$, then

$$V_g(\theta) \leq \pi_g^B(\theta) \leq 0,$$

so employing nobody weakly dominates employing only gender g.

Consequently, for each gender, either the benchmark contract satisfies the other gender's no-entry constraint and $V_g(\theta) = \pi_g^B(\theta)$, or the single-gender option is dominated by employing both genders or by employing nobody. We can therefore write

$$\begin{aligned}\Pi^*(\theta) &= \max\left\{V(\theta),\, V_M(\theta),\, V_F(\theta),\, 0\right\} \\ &= \max\left\{V(\theta),\, \pi_m^B(\theta),\, \pi_f^B(\theta),\, 0\right\}.\end{aligned}$$

It also follows that, whenever employing only one gender is strictly optimal, the relevant benchmark contract satisfies the other gender's no-entry constraint. The firm can therefore employ that gender with its efficient benchmark contract, without introducing any distortion.

B Proof of Lemma 2

To establish the chain of inequalities in (16), we rearrange $IC_f(\theta)$ and $IC_m(\theta)$ to obtain:

$$U_f(\theta) \geq U_m(\theta) - \gamma_f e_m(\theta) + \gamma_m e_m(\theta) = U_m(\theta) - \Delta\gamma e_m(\theta); \tag{A1}$$

$$U_m(\theta) \geq U_f(\theta) - \gamma_m e_f(\theta) + \gamma_f e_f(\theta) = U_f(\theta) + \Delta\gamma e_f(\theta). \tag{A2}$$

Equivalently,

$$IC_f(\theta) \text{ is satisfied if and only if } U_f(\theta) \geq U_m(\theta) - \Delta\gamma e_m(\theta); \tag{A3}$$

$$IC_m(\theta) \text{ is satisfied if and only if } U_m(\theta) \geq U_f(\theta) + \Delta\gamma e_f(\theta). \tag{A4}$$

Consequently, incentive compatibility requires:

$$U_m(\theta) - \Delta\gamma e_m(\theta) \leq U_f(\theta) \leq U_m(\theta) - \Delta\gamma e_f(\theta), \tag{A5}$$

The chain (A5) requires $U_m(\theta) - \Delta\gamma e_m(\theta) \leq U_m(\theta) - \Delta\gamma e_f(\theta)$, that is, $\Delta\gamma\left(e_m(\theta) - e_f(\theta)\right) \geq 0$. Since $\Delta\gamma > 0$, this gives $e_m(\theta) \geq e_f(\theta)$. Moreover, since $e_f(\theta) \geq 0$, the right-hand inequality in (A5) gives $U_f(\theta) \leq U_m(\theta) - \Delta\gamma e_f(\theta) \leq U_m(\theta)$.

C Proof of Proposition 2

To establish point (i) in the proposition, recall that Benchmark contracts are characterized by $U_m(\theta) = \overline{U}_m \equiv \overline{U}_m^k$, $U_f(\theta) = \overline{U}_f$ and first-best labor supplies, $e_g^B(\theta) = (1-\gamma_g)/\theta$, $g \in \{f, m\}$. Using (A3), Benchmark contracts satisfy $IC_f(\theta)$ if $\overline{U}_f \geq \overline{U}_m^k - \Delta\gamma e_m^B(\theta)$, or if:

$$\overline{U}_f \geq \overline{U}_m^k - \frac{(\gamma_f - \gamma_m)(1-\gamma_m)}{\theta}. \tag{A6}$$

In the same way, from (A4), Benchmark contracts satisfy $IC_m(\theta)$ if $\overline{U}_m^k \geq \overline{U}_f + \Delta\gamma e_f^B(\theta)$, or if:

$$\overline{U}_f \leq \overline{U}_m^k - \frac{(\gamma_f - \gamma_m)(1-\gamma_f)}{\theta}. \tag{A7}$$

Turning to point (ii), from (A6), solving for θ, $IC_f(\theta)$ is met for:

$$\theta \leq \theta_{IC_f}^B\left(\overline{U}_f\right) \equiv \frac{(\gamma_f - \gamma_m)(1-\gamma_m)}{\overline{U}_m^k - \overline{U}_f}. \tag{A8}$$

Similarly, from (A7), $IC_m(\theta)$ is met for:

$$\theta \geq \theta_{IC_m}^B\left(\overline{U}_f\right) \equiv \frac{(\gamma_f - \gamma_m)(1-\gamma_f)}{\overline{U}_m^k - \overline{U}_f}. \tag{A9}$$

Combining the last two conditions, Benchmark contracts with $e_g(\theta) = e_g^B(\theta)$ and $U_g(\theta) = \overline{U}_g$, $g \in \{f, m\}$, are incentive compatible for

$$\theta_{IC_m}^B\left(\overline{U}_f\right) \equiv \frac{(\gamma_f - \gamma_m)(1-\gamma_f)}{\overline{U}_m^k - \overline{U}_f} \leq \theta \leq \theta_{IC_f}^B\left(\overline{U}_f\right) \equiv \frac{(\gamma_f - \gamma_m)(1-\gamma_m)}{\overline{U}_m^k - \overline{U}_f}. \tag{A10}$$

Note that (A6) and (A7) are also the conditions under which the efficient contract of one gender satisfies the no-entry constraint of the other gender: women reject the efficient contract of men if and only if $\theta \leq \theta_{IC_f}^B(\overline{U}_f)$, and men reject the efficient contract of women if and only if $\theta \geq \theta_{IC_m}^B(\overline{U}_f)$.

Point (iii). Point I in Figure 1 is the intersection of the two benchmark employment cutoffs. Equating (8) and (9) gives

$$\frac{(1-\gamma_f)^2}{2\overline{U}_f} = \frac{(1-\gamma_m)^2}{2\overline{U}_m^k},$$

and therefore

$$\overline{U}_f^I = \overline{U}_m^k \frac{(1-\gamma_f)^2}{(1-\gamma_m)^2}, \qquad \theta^I = \tilde{\theta}_m^B = \frac{(1-\gamma_m)^2}{2\overline{U}_m^k}.$$

Substituting $\overline{U}_f^I$ into the two incentive-compatibility frontiers gives

$$\theta_{IC_m}^B(\overline{U}_f^I) = \frac{(1-\gamma_f)(1-\gamma_m)^2}{\overline{U}_m^k(2-\gamma_f-\gamma_m)} < \frac{(1-\gamma_m)^2}{2\overline{U}_m^k} < \frac{(1-\gamma_m)^3}{\overline{U}_m^k(2-\gamma_f-\gamma_m)} = \theta_{IC_f}^B(\overline{U}_f^I),$$

where both inequalities follow from $\gamma_f > \gamma_m$. Hence point I lies strictly inside the region in which the benchmark contracts are incentive compatible.

The remaining labeled points in Figure 1 are obtained in the same way, by equating a benchmark cutoff with an incentive-compatibility frontier. Specifically, A and C are the intersections of $\tilde{\theta}_m^B$ with $\theta_{IC_f}^B$ and $\theta_{IC_m}^B$, respectively, whereas H and K are the intersections of $\tilde{\theta}_f^B$ with $\theta_{IC_f}^B$ and $\theta_{IC_m}^B$, respectively. Solving these four equalities gives

$$
\begin{aligned}
A &= \left(\overline{U}_m^k \frac{1-2\gamma_f+\gamma_m}{1-\gamma_m}, \frac{(1-\gamma_m)^2}{2\overline{U}_m^k} \right), \\
H &= \left(\overline{U}_m^k \frac{(1-\gamma_f)^2}{(1-\gamma_f)^2+2\Delta\gamma(1-\gamma_m)}, \frac{(1-\gamma_f)^2+2\Delta\gamma(1-\gamma_m)}{2\overline{U}_m^k} \right), \\
C &= \left(\overline{U}_m^k \frac{(1-\gamma_f)^2+(\Delta\gamma)^2}{(1-\gamma_m)^2}, \frac{(1-\gamma_m)^2}{2\overline{U}_m^k} \right), \\
K &= \left(\overline{U}_m^k \frac{1-\gamma_f}{1+\gamma_f-2\gamma_m}, \frac{(1-\gamma_f)(1+\gamma_f-2\gamma_m)}{2\overline{U}_m^k} \right).
\end{aligned}
$$

All denominators are positive. Direct cross-multiplication gives the ordering of the horizontal coordinates

$$
\overline{U}_f^A < \overline{U}_f^H < \overline{U}_f^I < \overline{U}_f^K < \overline{U}_f^C. \tag{A11}
$$

The four successive comparisons reduce, respectively, to $2(\Delta\gamma)^3 > 0$, $(\Delta\gamma)^2 > 0$, $(\Delta\gamma)^2 > 0$, and $2(\Delta\gamma)^3 > 0$, after multiplication by the positive denominators. In particular, $\overline{U}_f^A < \overline{U}_f^I < \overline{U}_f^C$, which provides an alternative graphical verification that the two benchmark cutoffs intersect inside the incentive-compatible Benchmark region. The ordering $\overline{U}_f^H < \overline{U}_f^I < \overline{U}_f^K$ also shows how the female cutoff crosses the same region. These intersections delimit the single-gender regions AIH and ICK described in the main text.

D Proof of Proposition 3

Using U_f and U_m as choice variables instead of wages, problem (13) reads

$$
\begin{aligned}
\max_{e_f, e_m, U_f, U_m} \quad & \Phi \equiv S_f(e_f,\theta) + S_m(e_m,\theta) - U_f - U_m \\
\text{s.t.} \quad & U_f - \overline{U}_f \geq 0 \;\; [\lambda_f], \qquad U_m - \overline{U}_m^k \geq 0 \;\; [\lambda_m], \\
& U_f - U_m + \Delta\gamma e_m \geq 0 \;\; [\mu_f], \qquad U_m - U_f - \Delta\gamma e_f \geq 0 \;\; [\mu_m], \\
& e_f \geq 0, \qquad e_m \geq 0,
\end{aligned}
$$

where λ_f and λ_m are the multipliers associated with the participation constraints and μ_f and μ_m the ones associated with the incentive constraints. The objective is concave (strictly concave in efforts) and all constraints are linear. Because the constraints are linear, a constraint qualification holds, so the Kuhn–Tucker conditions are necessary at any optimum. Because the objective is concave, they are also sufficient for a global maximum. Finally, strict concavity in efforts implies that optimal effort levels are unique. The candidate effort levels below are strictly

positive on the interior of their respective regimes. Since the resulting efforts are positive, their non-negativity constraints do not affect the solution.

The Lagrangian is

$$\begin{aligned}\mathcal{L} = {} & S_f(e_f,\theta) + S_m(e_m,\theta) - U_f - U_m \\ & + \lambda_f(U_f - \overline{U}_f) + \lambda_m(U_m - \overline{U}_m^k) \\ & + \mu_f\big(U_f - U_m + \Delta\gamma e_m\big) + \mu_m\big(U_m - U_f - \Delta\gamma e_f\big).\end{aligned} \tag{A12}$$

Differentiating $\mathcal{L}$ with respect to the four choice variables gives the stationarity conditions

$$\partial\mathcal{L}/\partial U_f: \ -1 + \lambda_f + \mu_f - \mu_m = 0, \tag{A13}$$

$$\partial\mathcal{L}/\partial U_m: \ -1 + \lambda_m - \mu_f + \mu_m = 0, \tag{A14}$$

$$\partial\mathcal{L}/\partial e_f: \ 1 - \gamma_f - \theta e_f - \Delta\gamma\mu_m = 0 \iff e_f = e_f^B(\theta) - \frac{\Delta\gamma}{\theta}\mu_m, \tag{A15}$$

$$\partial\mathcal{L}/\partial e_m: \ 1 - \gamma_m - \theta e_m + \Delta\gamma\mu_f = 0 \iff e_m = e_m^B(\theta) + \frac{\Delta\gamma}{\theta}\mu_f. \tag{A16}$$

Conditions (A15) and (A16) show that the shadow price of men's incentive constraint, μ_m, reduces women's effort, whereas the shadow price of women's incentive constraint, μ_f, raises men's effort. We now show that, in each regime, the proposed menu and multipliers satisfy all Kuhn–Tucker conditions: (i) the stationarity conditions (A13)–(A16); (ii) non-negativity of the multipliers; (iii) primal feasibility, that is, the menu satisfies every constraint of problem (13); and (iv) complementary slackness, that is, $\lambda_f(U_f - \overline{U}_f) = 0$, $\lambda_m(U_m - \overline{U}_m^k) = 0$, $\mu_f(U_f - U_m + \Delta\gamma e_m) = 0$ and $\mu_m(U_m - U_f - \Delta\gamma e_f) = 0$, so that a constraint with a positive multiplier binds.

Regime *R*1. Set $\mu_f = \lambda_m = 0$. Then (A14) gives $\mu_m = 1$ and (A13) gives $\lambda_f = 2$. By (A15) and (A16), $e_f = e_f^{R1}(\theta)$ and $e_m = e_m^B(\theta)$. The binding constraints give $U_f = \overline{U}_f$ and $U_m = \overline{U}_f + \Delta\gamma e_f^{R1}(\theta)$. Hence, the multipliers are all non-negative, and stationarity and complementary slackness hold.

The slack constraint $PC_m(\theta)$ holds if and only if $\Delta\gamma e_f^{R1}(\theta) \geq \Delta\overline{U}$, that is, $\theta \leq \theta_{PC_m}(\overline{U}_f)$. The slack constraint $IC_f(\theta)$ holds because $U_f - U_m + \Delta\gamma e_m = \Delta\gamma\left(e_m^B(\theta) - e_f^{R1}(\theta)\right) = 2(\Delta\gamma)^2/\theta > 0$. Effort $e_f^{R1}(\theta)$ is positive if and only if Assumption 1 holds, which is also the condition under which $\theta_{PC_m}(\overline{U}_f) > 0$. This proves primal feasibility and identifies the frontier $\theta_{PC_m}(\overline{U}_f)$ of Figure 2 and expression (19), left part.

Regime *R*2. Set $\mu_f = 0$. Then (A16) gives $e_m = e_m^B(\theta)$. The binding constraints $PC_f(\theta)$, $PC_m(\theta)$, and $IC_m(\theta)$ give $U_f = \overline{U}_f$, $U_m = \overline{U}_m^k$, and $e_f = \Delta\overline{U}/\Delta\gamma = e_f^{R2}$. From (A15),

$$\mu_m = \frac{1}{\Delta\gamma}\left(1 - \gamma_f - \theta\frac{\Delta\overline{U}}{\Delta\gamma}\right),$$

and from (A13) and (A14), $\lambda_f = 1 + \mu_m$ and $\lambda_m = 1 - \mu_m$. Hence, stationarity holds. Complementary slackness also holds, because the only constraint that may be slack, $IC_f(\theta)$, has multiplier $\mu_f = 0$. The multipliers are non-negative if and only if $\theta_{PC_m}(\overline{U}_f) \leq \theta \leq \theta_{IC_m}^B(\overline{U}_f)$. Indeed,

$\mu_m \geq 0$ if and only if $\theta\Delta\overline{U} \leq \Delta\gamma(1-\gamma_f)$, that is, $\theta \leq \theta^B_{IC_m}(\overline{U}_f)$. Then $\lambda_f = 1+\mu_m \geq 1 > 0$. Finally, $\lambda_m \geq 0$ if and only if $\mu_m \leq 1$, that is, $\theta\Delta\overline{U}/\Delta\gamma \geq 1-\gamma_f-\Delta\gamma = 1-2\gamma_f+\gamma_m$, or $\theta \geq \theta_{PC_m}(\overline{U}_f)$.

The binding constraints hold with equality by construction. The slack constraint $IC_f(\theta)$ holds if and only if $\overline{U}_f - \overline{U}^k_m + \Delta\gamma e^B_m(\theta) \geq 0$, that is, $\theta \leq \theta^B_{IC_f}(\overline{U}_f)$. This is implied by $\theta \leq \theta^B_{IC_m}(\overline{U}_f)$, since $\theta^B_{IC_m}(\overline{U}_f) < \theta^B_{IC_f}(\overline{U}_f)$ when $\gamma_f > \gamma_m$. Both efforts are strictly positive: $e^B_m(\theta) > 0$, and $e^{R2}_f > 0$ because $\overline{U}_f < \overline{U}^k_m$. This proves primal feasibility. Unlike in Regime $R1$, the frontiers of Regime $R2$ come from the non-negativity of the multipliers rather than from primal feasibility: $\mu_m \geq 0$ gives $\theta^B_{IC_m}(\overline{U}_f)$, and $\lambda_m \geq 0$ gives $\theta_{PC_m}(\overline{U}_f)$, the same frontier found in Regime $R1$. At $\theta = \theta_{PC_m}(\overline{U}_f)$, the multipliers of the two regimes coincide ($\mu_m = 1$, $\lambda_m = 0$), so Regimes $R1$ and $R2$ join continuously.

Finally, we verify the ordering in (22). Since $e^{R1}_f(\theta) = (1-2\gamma_f+\gamma_m)/\theta$, $e^B_f(\theta) = (1-\gamma_f)/\theta$, and $\Delta\gamma, \Delta\overline{U} > 0$,

$$e^{R1}_f(\theta) \leq e^{R2}_f \iff \theta\,\Delta\overline{U} \geq \Delta\gamma\,(1-2\gamma_f+\gamma_m) \iff \theta \geq \theta_{PC_m}(\overline{U}_f),$$

$$e^{R2}_f \leq e^B_f(\theta) \iff \theta\,\Delta\overline{U} \leq \Delta\gamma\,(1-\gamma_f) \iff \theta \leq \theta^B_{IC_m}(\overline{U}_f).$$

Hence $e^{R1}_f(\theta) \leq e^{R2}_f \leq e^B_f(\theta)$ on Regime $R2$, with equality at the left and right frontier, respectively.

Regime B. Set $\mu_f = \mu_m = 0$ and $\lambda_f = \lambda_m = 1$. Effort levels are efficient and $U_g = \overline{U}_g$. The incentive constraints hold if and only if condition (A10) holds.

Regimes $R4$ and $R5$ are the mirror images of Regimes $R2$ and $R1$, respectively, with the roles of women and men exchanged. The Kuhn–Tucker conditions are verified in the same order: stationarity, complementary slackness, non-negativity of the multipliers, and primal feasibility. We therefore report only the main steps.

Regime $R4$. Set $\mu_m = 0$, so that $e_f = e^B_f(\theta)$. The binding constraints $PC_f(\theta)$, $PC_m(\theta)$, and $IC_f(\theta)$ give $U_f = \overline{U}_f$, $U_m = \overline{U}^k_m$, and $e_m = \Delta\overline{U}/\Delta\gamma = e^{R4}_m$. From (A16),

$$\mu_f = \frac{1}{\Delta\gamma}\left(\theta\frac{\Delta\overline{U}}{\Delta\gamma} - (1-\gamma_m)\right),$$

which is non-negative if and only if $\theta \geq \theta^B_{IC_f}(\overline{U}_f)$. From (A13) and (A14), $\lambda_m = 1+\mu_f \geq 0$ and $\lambda_f = 1-\mu_f$, which is non-negative if and only if $\theta \leq \theta_{PC_f}(\overline{U}_f)$. The slack constraint $IC_m(\theta)$ holds if and only if $\theta \geq \theta^B_{IC_m}(\overline{U}_f)$, which is implied by $\theta \geq \theta^B_{IC_f}(\overline{U}_f)$.

Finally, we verify the ordering in (23). Since $e^B_m(\theta) = (1-\gamma_m)/\theta$, $e^{R5}_m(\theta) = (1-2\gamma_m+\gamma_f)/\theta$, and $\Delta\gamma, \Delta\overline{U} > 0$,

$$e^B_m(\theta) \leq e^{R4}_m \iff \theta\,\Delta\overline{U} \geq \Delta\gamma\,(1-\gamma_m) \iff \theta \geq \theta^B_{IC_f}(\overline{U}_f),$$

$$e^{R4}_m \leq e^{R5}_m(\theta) \iff \theta\,\Delta\overline{U} \leq \Delta\gamma\,(1-2\gamma_m+\gamma_f) \iff \theta \leq \theta_{PC_f}(\overline{U}_f).$$

Hence $e_m^B(\theta) \leq e_m^{R4} \leq e_m^{R5}(\theta)$ on Regime $R4$, with equality at the left and right frontier, respectively.

Regime $R5$. Set $\mu_m = \lambda_f = 0$. Then (A13) gives $\mu_f = 1$ and (A14) gives $\lambda_m = 2$. By (A15) and (A16), $e_f = e_f^B(\theta)$ and $e_m = e_m^{R5}(\theta)$. The binding constraints give $U_m = \overline{U}_m^k$ and $U_f = \overline{U}_m^k - \Delta\gamma e_m^{R5}(\theta)$. The slack constraint $PC_f(\theta)$ holds if and only if $\Delta\overline{U} \geq \Delta\gamma e_m^{R5}(\theta)$, that is, $\theta \geq \theta_{PC_f}(\overline{U}_f)$. The slack constraint $IC_m(\theta)$ holds because $U_m - U_f - \Delta\gamma e_f = \Delta\gamma \left(e_m^{R5}(\theta) - e_f^B(\theta)\right) = 2(\Delta\gamma)^2/\theta > 0$.

By (20), the five intervals of θ cover $(0, \infty)$. At each frontier, the solutions of the two adjacent regimes coincide (for instance, at $\theta = \theta_{PC_m}(\overline{U}_f)$, $e_f^{R1}(\theta) = e_f^{R2}$ and men's rent in $R1$ is zero), so that effort levels and utilities are continuous in θ and $\overline{U}_f$. Finally, expression (25) follows from $S_g(e_g, \theta) - U_g = \pi_g^B(\theta) - \frac{1}{2}\theta(e_g - e_g^B(\theta))^2 - (U_g - \overline{U}_g)$, which is a direct consequence of (11) and (12).

E Proof of Lemma 3 and Proposition 4

E.1 Proof of Lemma 3

Structure of the proof. Points (i) and (ii) apply the same argument to the two cases in which the female and male contracts, respectively, are mimicked. First, we substitute the menus characterized in Proposition 3 into the definition of the relevant net contribution. This expresses the net contribution as the corresponding benchmark profit minus the cost of the effort distortion and any rent paid to the mimicking gender. It follows that the net contribution cannot exceed the benchmark profit and coincides with it at the relevant incentive-compatibility frontier, where these incentive costs disappear. The explicit formulas then establish strict monotonicity in θ; continuity follows because the menus of the two adjacent distorted regimes coincide at their common frontier. Finally, the location of each distorted region relative to the incentive-compatibility frontier implies that a positive net contribution from the gender whose contract is mimicked is possible only when the benchmark profit from the other gender is also positive. This last property is what allows the hiring rule in Proposition 4 to be stated in terms of the sign of the employment contributions.

Point (i). In Regimes $R1$ and $R2$, $e_m = e_m^B(\theta)$ and $U_f = \overline{U}_f$, so that (25) gives $V(\theta) = \pi_m^B(\theta) + \pi_f^{GN}(\theta)$ with $\pi_f^{GN}(\theta)$ defined in (26). In $R1$, $e_f^{R1}(\theta) - e_f^B(\theta) = -\Delta\gamma/\theta$ and $U_m - \overline{U}_m^k = \Delta\gamma e_f^{R1}(\theta) - \Delta\overline{U}$, so that

$$\pi_f^{GN}(\theta) = \frac{(1-\gamma_f)^2}{2\theta} - \overline{U}_f - \frac{(\Delta\gamma)^2}{2\theta} - \frac{\Delta\gamma(1-\gamma_f-\Delta\gamma)}{\theta} + \Delta\overline{U} = \frac{(1-\gamma_f-\Delta\gamma)^2}{2\theta} - \left(2\overline{U}_f - \overline{U}_m^k\right),$$

which is the first line of (27) since $1 - \gamma_f - \Delta\gamma = 1 - 2\gamma_f + \gamma_m$. In $R2$, men receive no rent, so that $\pi_f^{GN}(\theta) = S_f(e_f^{R2}, \theta) - \overline{U}_f$, which is the second line of (27).

By (26), $\pi_f^{GN}(\theta) \leq \pi_f^B(\theta)$, with equality if and only if $e_f = e_f^B(\theta)$ and $U_m = \overline{U}_m^k$. In $R1$,

$e_f^{R1}(\theta) < e_f^B(\theta)$; in $R2$, $e_f^{R2} = e_f^B(\theta)$ if and only if $\theta = \theta^B_{IC_m}(\overline{U}_f)$. Hence equality holds only at $\theta = \theta^B_{IC_m}(\overline{U}_f)$.

$\pi_f^{GN}(\theta)$ is strictly decreasing in θ: its derivative is $-(1-2\gamma_f+\gamma_m)^2/(2\theta^2) < 0$ in $R1$ and $-\frac{1}{2}(\Delta\overline{U}/\Delta\gamma)^2 < 0$ in $R2$ (recall that $\overline{U}_f < \overline{U}^k_m$). It is continuous at $\theta_{PC_m}(\overline{U}_f)$, because the menus of $R1$ and $R2$ coincide there.

The derivative of $\pi_f^{GN}(\theta)$ with respect to $\overline{U}_f$ follows from the envelope theorem. $\overline{U}_f$ enters problem (13) only through $PC_f(\theta)$; hence, from the Lagrangian (A12), we have $\partial V(\theta)/\partial\overline{U}_f = -\lambda_f$, and (A13) with $\mu_f = 0$ gives $\lambda_f = 1+\mu_m$. Since $V(\theta) = \pi_m^B(\theta) + \pi_f^{GN}(\theta)$ and $\pi_m^B(\theta)$ does not depend on $\overline{U}_f$, it follows that $\partial\pi_f^{GN}(\theta)/\partial\overline{U}_f = -(1+\mu_m)$. Direct differentiation of (27) confirms this. In Regime $R1$, the first line of (27) gives

$$\frac{\partial\pi_f^{GN}(\theta)}{\partial\overline{U}_f} = -2 = -(1+\mu_m), \qquad \text{since } \mu_m = 1.$$

In Regime $R2$, the second line of (27), together with $\partial\Delta\overline{U}/\partial\overline{U}_f = -1$, gives

$$\frac{\partial\pi_f^{GN}(\theta)}{\partial\overline{U}_f} = -\frac{1-\gamma_f}{\Delta\gamma} + \theta\frac{\Delta\overline{U}}{(\Delta\gamma)^2} - 1 = -(1+\mu_m),$$

Since $\mu_m \in [0,1]$ in $R2$ (see Appendix D), this derivative lies in $[-2,-1]$: it equals -1, the derivative of $\pi_f^B(\theta)$, at $\theta = \theta^B_{IC_m}(\overline{U}_f)$, and -2, its value in $R1$, at $\theta = \theta_{PC_m}(\overline{U}_f)$.

Finally, we show that $\pi_f^{GN}(\theta) > 0$ implies $\pi_m^B(\theta) > 0$. Suppose that $\pi_f^{GN}(\theta) > 0$. Then $\pi_f^B(\theta) > 0$, that is, $2\overline{U}_f\theta < (1-\gamma_f)^2$. Moreover, $\theta \le \theta^B_{IC_m}(\overline{U}_f)$ means $\theta\Delta\overline{U} \le \Delta\gamma(1-\gamma_f)$. Adding twice the second inequality to the first,

$$2\overline{U}^k_m\theta < (1-\gamma_f)^2 + 2\Delta\gamma(1-\gamma_f) = (1-\gamma_m)^2 - (\Delta\gamma)^2 < (1-\gamma_m)^2,$$

so that $\pi_m^B(\theta) > 0$.

Point (ii). The argument is symmetric. In $R4$ and $R5$, $e_f = e_f^B(\theta)$ and $U_m = \overline{U}^k_m$, so that $V(\theta) = \pi_f^B(\theta) + \pi_m^{GN}(\theta)$. In $R5$, $e_m^{R5}(\theta) - e_m^B(\theta) = \Delta\gamma/\theta$ and $U_f - \overline{U}_f = \Delta\overline{U} - \Delta\gamma e_m^{R5}(\theta)$, so that

$$\pi_m^{GN}(\theta) = \frac{(1-\gamma_m)^2}{2\theta} - \overline{U}^k_m - \frac{(\Delta\gamma)^2}{2\theta} + \frac{\Delta\gamma(1-\gamma_m+\Delta\gamma)}{\theta} - \Delta\overline{U} = \frac{(1-\gamma_m+\Delta\gamma)^2}{2\theta} - \left(2\overline{U}^k_m - \overline{U}_f\right).$$

In $R4$, $\pi_m^{GN}(\theta) = S_m(e_m^{R4},\theta) - \overline{U}^k_m$. The inequality $\pi_m^{GN}(\theta) \le \pi_m^B(\theta)$, with equality only at $\theta^B_{IC_f}(\overline{U}_f)$, and the strict monotonicity in θ follow as in point (i). By the same envelope argument, now with $\mu_m = 0$, (A13) gives $\lambda_f = 1-\mu_f$. Since $V(\theta) = \pi_f^B(\theta) + \pi_m^{GN}(\theta)$ and $\partial\pi_f^B(\theta)/\partial\overline{U}_f = -1$, we have $\partial\pi_m^{GN}(\theta)/\partial\overline{U}_f = -\lambda_f + 1 = \mu_f$. Direct differentiation of (29) confirms this: it gives $1 = \mu_f$ in $R5$ and $-(1-\gamma_m)/\Delta\gamma + \theta\Delta\overline{U}/(\Delta\gamma)^2 = \mu_f$ in $R4$.

Finally, if $\pi_m^{GN}(\theta) > 0$, then $2\overline{U}^k_m\theta < (1-\gamma_m)^2$; moreover, $\theta \ge \theta^B_{IC_f}(\overline{U}_f)$ means $\theta\Delta\overline{U} \ge \Delta\gamma(1-\gamma_m)$. Subtracting twice the second inequality from the first,

$$2\overline{U}_f\theta < (1-\gamma_m)^2 - 2\Delta\gamma(1-\gamma_m) = (1-\gamma_f)^2 - (\Delta\gamma)^2 < (1-\gamma_f)^2,$$

so that $\pi_f^B(\theta) > 0$.

E.2 Proof of Proposition 4

Structure of the proof. In each distorted regime, let g denote the gender whose contract is mimicked and let $g' \neq g$ denote the other gender. Thus, $(g, g') = (f, m)$ in $R1$–$R2$, while $(g, g') = (m, f)$ in $R4$–$R5$. The contract of g' remains efficient, so the proof takes employment of g' alone as the baseline and asks whether adding g raises the firm's profit. The gain from adding g is exactly the net contribution of g; hence, its sign determines whether the firm employs one or both genders. Lemmas 1 and 3 ensure that no other employment choice can yield a higher profit. The proof then considers the Benchmark region, where there is no incentive problem and the hiring decision for each gender depends directly on its benchmark profit. The two arguments coincide at the incentive-compatibility frontiers and together establish the pointwise hiring rule.

Write π_g^{GN} for the relevant net contribution, namely π_f^{GN} when $g = f$ and π_m^{GN} when $g = m$. In either pair of distorted regimes,

$$V(\theta) = \pi_{g'}^B(\theta) + \pi_g^{GN}(\theta).$$

In the interior of either pair of distorted regimes, employing g alone is never optimal. If $\pi_g^B(\theta) \leq 0$, it is weakly dominated by employing nobody. If $\pi_g^B(\theta) > 0$, the benchmark contract of g strictly attracts g'; by Lemma 1(i), offering that contract to both genders yields at least

$$2\pi_g^B(\theta) > \pi_g^B(\theta).$$

The firm therefore compares only

$$\pi_{g'}^B(\theta) + \pi_g^{GN}(\theta), \qquad \pi_{g'}^B(\theta), \qquad 0.$$

The first option adds $\pi_g^{GN}(\theta)$ to the second, and the second adds $\pi_{g'}^B(\theta)$ to the third. Moreover, Lemma 3 implies

$$\pi_g^{GN}(\theta) > 0 \quad \Longrightarrow \quad \pi_{g'}^B(\theta) > 0.$$

Consequently, both genders are employed when $\pi_g^{GN}(\theta) > 0$; only g' is employed when $\pi_g^{GN}(\theta) \leq 0 < \pi_{g'}^B(\theta)$; and nobody is employed when $\pi_{g'}^B(\theta) \leq 0$. This is precisely the rule $\pi_g^{\text{emp}}(\theta) > 0$ and $\pi_{g'}^{\text{emp}}(\theta) > 0$ in (31).

In Regime B, including its boundaries, the benchmark contracts are incentive compatible: neither gender strictly prefers the other gender's contract. Hence the problem separates:

$$\Pi^*(\theta) = \max\{\pi_f^B(\theta), 0\} + \max\{\pi_m^B(\theta), 0\},$$

which again gives (31). At the two incentive-compatibility frontiers, the net contribution equals the corresponding benchmark profit, so the rules agree continuously.

Finally, Lemma 1(ii) implies that a gender employed alone receives its efficient contract at its outside option. The no-entry constraint makes the other gender weakly prefer its outside option; our tie-breaking convention assigns that option in case of indifference.

F Proof of Proposition 5 and Corollary 1

F.1 Coordinates and ordering of the relevant points

The coordinates of points A, H, I, K, and C, which describe the intersections among the benchmark cutoffs and the incentive-compatibility frontiers in Figure 1, are derived in Appendix C. It remains to derive the two points that involve the distorted employment cutoffs in Figure 2. Their coordinates $(\overline{U}_f, \theta)$ are

$$F = \left(\overline{U}_m^k \frac{1-\gamma_m}{1+2\gamma_f-3\gamma_m}, \frac{(1-2\gamma_f+\gamma_m)(1+2\gamma_f-3\gamma_m)}{2\overline{U}_m^k} \right),$$
$$Z = \left(\overline{U}_m^k \frac{1-3\gamma_f+2\gamma_m}{1-\gamma_f}, \frac{(1-\gamma_f)(1+\gamma_f-2\gamma_m)}{2\overline{U}_m^k} \right).$$

Points F and Z are the points at which the distorted cutoffs cross the frontiers $\theta_{PC_m}(\overline{U}_f)$ and $\theta_{PC_f}(\overline{U}_f)$, respectively. Point F is defined by $\pi_f^{GN}\left(\theta_{PC_m}(\overline{U}_f)\right) = 0$: women's net contribution is zero exactly at the boundary between Regimes $R1$ and $R2$. Evaluating the first line of (27) at $\theta_{PC_m}(\overline{U}_f) = \Delta\gamma(1-2\gamma_f+\gamma_m)/\Delta\overline{U}$ gives

$$\pi_f^{GN}\left(\theta_{PC_m}(\overline{U}_f)\right) = \frac{(1-2\gamma_f+\gamma_m)\Delta\overline{U}}{2\Delta\gamma} - \left(2\overline{U}_f - \overline{U}_m^k\right).$$

Setting this expression equal to zero and using $\Delta\overline{U} = \overline{U}_m^k - \overline{U}_f$ gives

$$(1-2\gamma_f+\gamma_m)\left(\overline{U}_m^k - \overline{U}_f\right) = 2\Delta\gamma\left(2\overline{U}_f - \overline{U}_m^k\right).$$

Since $1-2\gamma_f+\gamma_m+2\Delta\gamma = 1-\gamma_m$ and $1-2\gamma_f+\gamma_m+4\Delta\gamma = 1+2\gamma_f-3\gamma_m$, the solution is $\overline{U}_f^F = \overline{U}_m^k(1-\gamma_m)/(1+2\gamma_f-3\gamma_m)$. The θ-coordinate of F is $\theta_{PC_m}(\overline{U}_f^F)$. Since $\overline{U}_m^k - \overline{U}_f^F = 2\Delta\gamma\,\overline{U}_m^k/(1+2\gamma_f-3\gamma_m)$, it equals $(1-2\gamma_f+\gamma_m)(1+2\gamma_f-3\gamma_m)/(2\overline{U}_m^k)$.

Symmetrically, point Z is defined by $\pi_m^{GN}\left(\theta_{PC_f}(\overline{U}_f)\right) = 0$. Evaluating the second line of (29) at $\theta_{PC_f}(\overline{U}_f) = \Delta\gamma(1-2\gamma_m+\gamma_f)/\Delta\overline{U}$ gives

$$\pi_m^{GN}\left(\theta_{PC_f}(\overline{U}_f)\right) = \frac{(1-2\gamma_m+\gamma_f)\Delta\overline{U}}{2\Delta\gamma} - \left(2\overline{U}_m^k - \overline{U}_f\right).$$

Setting this expression equal to zero gives

$$(1-2\gamma_m+\gamma_f)\left(\overline{U}_m^k - \overline{U}_f\right) = 2\Delta\gamma\left(2\overline{U}_m^k - \overline{U}_f\right).$$

Since $1-2\gamma_m+\gamma_f-2\Delta\gamma = 1-\gamma_f$ and $1-2\gamma_m+\gamma_f-4\Delta\gamma = 1-3\gamma_f+2\gamma_m$, the solution is $\overline{U}_f^Z = \overline{U}_m^k(1-3\gamma_f+2\gamma_m)/(1-\gamma_f)$. The θ-coordinate of Z is $\theta_{PC_f}(\overline{U}_f^Z)$. Since $\overline{U}_m^k - \overline{U}_f^Z = 2\Delta\gamma\,\overline{U}_m^k/(1-\gamma_f)$ and $1-2\gamma_m+\gamma_f = 1+\gamma_f-2\gamma_m$, it equals $(1-\gamma_f)(1+\gamma_f-2\gamma_m)/(2\overline{U}_m^k)$.

All denominators used in these comparisons are positive, since $1-\gamma_f > 0$ and $\gamma_f \geq \gamma_m$. Appendix C establishes $\overline{U}_f^A < \overline{U}_f^H < \overline{U}_f^I < \overline{U}_f^K < \overline{U}_f^C$. The remaining comparisons follow by

direct cross-multiplication: $Z < A$ and $K < F$ each reduce to $2(\Delta\gamma)^2 > 0$, while $F < \overline{U}_m^k$ reduces to $2\Delta\gamma > 0$. Hence

$$\overline{U}_f^Z < \overline{U}_f^A < \overline{U}_f^H < \overline{U}_f^I < \overline{U}_f^K < \overline{U}_f^C, \qquad \overline{U}_f^K < \overline{U}_f^F < \overline{U}_m^k.$$

Assumption 1 is not needed for the ordering itself. It ensures that $\overline{U}_f^A > 0$ and that $2\overline{U}_f^F - \overline{U}_m^k > 0$ (see the proof of Proposition 5). Points K and F delimit the two pieces of the distorted female cutoff, while points Z and A delimit the two pieces of the distorted male cutoff.

F.2 Proof of Proposition 5

The proof uses the effective contributions in (30). By Lemma 3, $\pi_f^{\text{emp}}(\theta)$ and $\pi_m^{\text{emp}}(\theta)$ are continuous at their respective incentive-compatibility frontiers and strictly decreasing in θ. Moreover, under Assumption 1,

$$\lim_{\theta\to 0} \pi_f^{\text{emp}}(\theta) = +\infty, \qquad \lim_{\theta\to\infty} \pi_f^{\text{emp}}(\theta) = -\overline{U}_f < 0,$$

and

$$\lim_{\theta\to 0} \pi_m^{\text{emp}}(\theta) = +\infty, \qquad \lim_{\theta\to\infty} \pi_m^{\text{emp}}(\theta) = -\left(2\overline{U}_m^k - \overline{U}_f\right) < 0.$$

Hence, each employment contribution has exactly one finite zero, and (31) implies that this zero is the corresponding employment cutoff.

Female cutoff. Point K is defined by

$$\theta_{IC_m}^B(\overline{U}_f^K) = \tilde{\theta}_f^B(\overline{U}_f^K).$$

If $\overline{U}_f \le \overline{U}_f^K$, the frontier lies weakly below the benchmark cutoff. At the frontier, $\pi_f^{GN} = \pi_f^B \ge 0$; since π_f^{GN} is decreasing, it is positive below the frontier. Above the frontier, $\pi_f^{\text{emp}} = \pi_f^B$. Therefore the unique zero of π_f^{emp} is $\tilde{\theta}_f^B$.

If $\overline{U}_f > \overline{U}_f^K$, the benchmark cutoff lies below the frontier, so

$$\pi_f^{GN}\left(\theta_{IC_m}^B(\overline{U}_f)\right) = \pi_f^B\left(\theta_{IC_m}^B(\overline{U}_f)\right) < 0.$$

Because π_f^{GN} is strictly decreasing and tends to $+\infty$ as $\theta \to 0$, its unique zero lies below the frontier and is the female cutoff. To locate this zero, evaluate the net contribution at the boundary between $R1$ and $R2$:

$$\pi_f^{GN}\left(\theta_{PC_m}(\overline{U}_f)\right) = \frac{(1 - 2\gamma_f + \gamma_m)\Delta\overline{U}}{2\Delta\gamma} - \left(2\overline{U}_f - \overline{U}_m^k\right).$$

This expression is zero at F and strictly decreasing in $\overline{U}_f$. Hence the zero of π_f^{GN} lies in $R2$ between K and F, and in $R1$ beyond F. Setting the corresponding line of (27) equal to zero gives

$$\tilde{\theta}_f^{R2} = 2\frac{\Delta\gamma}{(\Delta\overline{U})^2}\left[\overline{U}_m^k(1-\gamma_f) - \overline{U}_f(1-\gamma_m)\right], \qquad \tilde{\theta}_f^{R1} = \frac{(1 - 2\gamma_f + \gamma_m)^2}{2(2\overline{U}_f - \overline{U}_m^k)}.$$

The formulas coincide at K and F, because at each point the same continuous contribution is zero. This proves point (i).

Male cutoff. The argument is symmetric. Point A is defined by

$$\theta^B_{IC_f}(\overline{U}^A_f) = \tilde{\theta}^B_m.$$

If $\overline{U}_f \geq \overline{U}^A_f$, the benchmark cutoff lies weakly below the frontier, and the unique zero of π^{emp}_m is $\tilde{\theta}^B_m$. If $\overline{U}_f < \overline{U}^A_f$, then $\pi^{GN}_m = \pi^B_m > 0$ at the frontier, while $\pi^{GN}_m < 0$ for sufficiently large θ. Its unique zero above the frontier is therefore the male cutoff.

At the boundary between $R4$ and $R5$,

$$\pi^{GN}_m\big(\theta_{PC_f}(\overline{U}_f)\big) = \frac{(1-2\gamma_m+\gamma_f)\Delta\overline{U}}{2\Delta\gamma} - \left(2\overline{U}^k_m - \overline{U}_f\right).$$

This expression is zero at Z and strictly decreasing in $\overline{U}_f$. Hence the root lies in $R5$ to the left of Z, and in $R4$ between Z and A. Solving $\pi^{GN}_m = 0$ gives

$$\tilde{\theta}^{R5}_m = \frac{(1-2\gamma_m+\gamma_f)^2}{2(2\overline{U}^k_m - \overline{U}_f)}, \qquad \tilde{\theta}^{R4}_m = 2\frac{\Delta\gamma}{(\Delta\overline{U})^2}\left[\overline{U}^k_m(1-\gamma_f) - \overline{U}_f(1-\gamma_m)\right].$$

Continuity at Z and A again follows because the same contribution is zero on both sides. This proves point (*ii*).

Additional exclusion and gender composition. Whenever a cutoff is distorted, it is the zero of $\pi^{GN}_g < \pi^B_g$. At that zero, benchmark profit is still positive; hence the gender-neutral cutoff lies strictly below the tagging cutoff. This happens for women exactly to the right of K, and for men exactly to the left of A, proving point (*iii*).

For point (*iv*), recall from Appendix C that

$$\overline{U}^A_f < \overline{U}^I_f < \overline{U}^K_f, \qquad \tilde{\theta}^B_f(\overline{U}^I_f) = \tilde{\theta}^B_m.$$

Between A and K, both cutoffs equal their benchmark values and cross at I. To the right of K, the female cutoff is weakly below its benchmark value while the male cutoff equals its benchmark value; therefore $\tilde{\theta}_f < \tilde{\theta}_m$. To the left of A, the symmetric argument gives $\tilde{\theta}_f > \tilde{\theta}_m$. Thus, when both cutoffs are inside the support, more men are employed if and only if $\overline{U}_f > \overline{U}^I_f$.

F.3 Proof of Corollary 1

For women, direct differentiation in $R1$ gives

$$\frac{d\tilde{\theta}^{R1}_f}{d\overline{U}_f} < 0, \qquad \frac{d}{d\overline{U}_f}\left(\frac{\tilde{\theta}^{R1}_f}{\tilde{\theta}^B_f}\right) = -\frac{(1-2\gamma_f+\gamma_m)^2}{(1-\gamma_f)^2}\frac{\overline{U}^k_m}{(2\overline{U}_f - \overline{U}^k_m)^2} < 0.$$

In $R2$,

$$\frac{d\tilde{\theta}^{R2}_f}{d\overline{U}_f} = \frac{2\Delta\gamma\left[\overline{U}^k_m(1-2\gamma_f+\gamma_m) - \overline{U}_f(1-\gamma_m)\right]}{(\Delta\overline{U})^3},$$

which is negative if and only if $\overline{U}_f > \overline{U}_f^A$, and hence throughout the relevant interval $[\overline{U}_f^K, \overline{U}_f^F]$. Moreover, the sign of

$$\frac{d}{d\overline{U}_f}\left(\frac{\tilde{\theta}_f^{R2}}{\tilde{\theta}_f^B}\right)$$

is the sign of

$$\overline{U}_m^k(1-\gamma_f) - \overline{U}_f(1+\gamma_f-2\gamma_m),$$

which is negative exactly when $\overline{U}_f > \overline{U}_f^K$. Since the ratio is continuous at F and equals one at K, point (i) follows.

For men, $\tilde{\theta}_m^{R5}$ is strictly increasing because its denominator $2\overline{U}_m^k - \overline{U}_f$ is positive and decreasing in $\overline{U}_f$. The formula for $\tilde{\theta}_m^{R4}$ is the same as that for $\tilde{\theta}_f^{R2}$; the derivative displayed above is positive if and only if $\overline{U}_f < \overline{U}_f^A$, and hence throughout the relevant $R4$ interval. The male cutoff is continuous at Z and A. Finally, $\tilde{\theta}_m^B$ is independent of $\overline{U}_f$, so the ratio $\tilde{\theta}_m/\tilde{\theta}_m^B$ has the same monotonicity as the male cutoff and reaches one at A. This proves point (ii).