Internal Pay Equity and the Quantity-Quality Trade-Off in Hiring*

Michael Amior† and Shmuel San‡

March 18, 2025

Abstract

Firms face significant constraints in their ability to differentiate pay by worker productivity. We show how these internal equity constraints generate a quantity-quality trade-off in hiring: firms which offer higher wages attract higher skilled workers, but cannot profitably employ lower skilled workers. In equilibrium, this mechanism leads to workplace segregation and pay dispersion even among ex-ante identical firms. Unlike in a conventional monopsony model, firms use higher pay to improve hiring quality, even at the cost of lower quantity. Our framework provides a novel interpretation of the (empirically successful) log additive AKM wage model, and shows how log additivity can be reconciled with sorting of high-skilled workers to high-paying firms. It can also rationalize a novel hump-shaped relationship between firm size and firm pay (which we document using Israeli data) and, by implication, the negligible wage return to firm size. Finally, our model provides new insights into aggregate-level and regional changes in worker-firm sorting and earnings inequality, firms' location choices, and public-private sector wage differentials—which we explore empirically.

*We are grateful to Momi Dahan, Oren Danieli and Pat Kline for helpful comments, and to seminar participants at Ben Gurion, Tel Aviv and Hebrew University. We also thank Matan Goldman for excellent research assistance, and David Gordon, Anat Katz, Rebecca Krieger, Julia Vidar, and many other employees of the Israeli CBS, for their invaluable administrative support. This research was generously supported by Israel Science Foundation grants 1277/21 (Amior) and 1472/22 (San), and the Maurice Falk Institute.

†Hebrew University of Jerusalem. Email: **michael.amior@mail.huji.ac.il**

‡Hebrew University of Jerusalem. Email: **muly.san@mail.huji.ac.il**.

1 Introduction

Firms face significant constraints in their ability to differentiate pay between workers, stemming from workers' equity concerns. These constraints manifest both horizontally—between workers performing similar jobs—and vertically—across different levels of a firm's hierarchy (Akerlof and Yellen, 1990; Manning, 1994; Bewley, 1999; Weil, 2014; Giupponi and Machin, 2022). Empirical studies from diverse contexts show that perceived pay inequity or unfairness can harm effort, group morale and retention (Card et al., 2012; Breza et al., 2018; Dube et al., 2019). In this paper, we explore how equity constraints influence firms' pay and recruitment strategies in labor market equilibrium. We argue that this can generate a quantity-quality trade-off in hiring, which sheds new light on numerous labor market phenomena.

Our point of departure is the monopsony model of Card et al. (2018), where firms' wage-setting power is predicated on workers' idiosyncratic preferences over workplaces. Within this setting, we impose a strict limit on the extent to which firms can differentiate pay by worker productivity (motivated by Akerlof and Yellen, 1990). When this constraint binds, firms must trade off quantity with quality in hiring: higher wages help attract higher-skilled workers, but make it unprofitable to employ lower-skilled workers. This trade-off sustains two distinct firm strategies in equilibrium: (i) a "selective" strategy, paying high wages to recruit high-skill workers, while rationing low-skill employment, and (ii) an "inclusive" strategy, paying lower wages to maintain a larger, more diverse workforce. The prevalence of selective firms (and hence aggregate earnings inequality) is increasing in both the bite of the equity constraint and the abundance (and productivity) of high-skilled labor.[1] This results in substantial workplace segregation and firm wage dispersion, even among ex-ante identical firms.

Our framework implies that wages are log additive in fixed firm and worker effects, in common with the "AKM" wage model (Abowd et al., 1999). Intuitively, a binding equity constraint compels firms to adopt a single proportional pay premium (or "company wage policy", in the language of Manning, 1994), which they apply uniformly to their workforce. Though the AKM model is often chosen for econometric convenience, it happens to fit the data remarkably well in numerous settings (Card et al., 2013; Kline, 2024); and our framework provides a novel conceptual interpretation.

In addition, as Bonhomme et al. (2019) and Kline (2025) emphasize, it is difficult to reconcile log additive wages with the observed heavy sorting of high-skilled workers to high-paying firms, if sorting is driven by complementarities in production (as in Becker, 1973).

[1] This positive effect of skill supply on earnings inequality is shared with the directed technical change model of Acemoglu (1998), but the story here is very different.

But positive sorting is a natural implication of our framework: if the equity constraint binds, firms will use higher pay to improve hiring *quality*, even at the cost of lower *quantity* (in stark contrast to conventional models), and even in the absence of productive complementarities between firms and workers. Interestingly, unlike in settings with productive complementarities, this sorting is socially inefficient: it reflects an inability of high-paying firms to employ low-skilled workers at wages commensurate with their productivity.

While the quantity-quality trade-off generates a positive relationship between firm pay and hiring quality, it also mutes its relationship with workforce size. Once we allow for (skill-neutral) variation in productivity across firms, our model implies a concave or even hump-shaped relationship between workforce size and firm pay. This is because the density of selective firms grows more quickly higher up the firm pay distribution, so the quantity-quality trade-off becomes more acute. This insight can help explain the surprisingly small wage return to log firm size in cross-sectional data, typically estimated at around 0.05 (Sokolova and Sorensen, 2021). Conventional monopsony models would require implausibly elastic labor supply to individual firms to generate such small premia (Bloesch and Larsen, 2023), but they are a natural consequence of binding equity constraints in our framework.[2]

We test the model's predictions using Israeli administrative data from 1990 to 2019, which provides detailed information on workers' education, wages, and employment histories. The Israeli context is particularly suitable for this analysis, as its well-documented tech boom provides valuable empirical variation. The period saw large growth in workforce education, particularly in STEM degrees—coinciding with a rapid increase in the wage returns to these degrees. Our core empirical analysis focuses on cross-sectional variation in wages, employment and skill shares across firms. But our model also makes predictions for market-level variation, as the prevalence of highly productive skill labor should affect the attractiveness of the selective hiring strategy. To this end, we exploit the substantial variation in skill shares across regions and over time, afforded by our setting.

The empirical evidence strongly supports our theoretical framework. First, we document that the relationship between firm size and wages follows an inverse-U shape, consistent with the quantity-quality trade-off in our model. This pattern is entirely attributable to low-educated workers, whereas employment of high-educated workers increases monotonically with firm wages—just as our model predicts. We also see the same patterns within detailed industry categories. To our knowledge, this hump-shape relationship is new to the literature, and it appears to be a general phenomenon: we document similar patterns in Northern Italy using the Veneto Worker History (VWH) file, a popular dataset in the labor literature.

[2]Bloesch and Larsen (2023) offer an alternative story for small firm size premia, arising from a recruitment expenditure margin which can shape hiring quantities independently of wages (as in Manning, 2006).

Despite the heavy sorting of high-skilled workers to high-paying firms, we show that firm wage premia are remarkably similar across education groups—consistent with the log additive AKM wage model and previous empirical work.

These qualitative patterns offer compelling support for our interpretation of the data. But we also fit the data to our very parsimonious model, using a three-group skill classification (non-graduates, non-STEM graduates, and STEM graduates) and skill-neutral firm heterogeneity. Despite its simplicity, our model is able to match the key results surprisingly well: (i) log additive wages, (ii) skill sorting patterns, and (iii) the hump-shaped size-pay relationship (and small firm size premium). Our estimates imply that the equity constraint compresses the STEM degree return by 31% within firms (relative to the productivity differential), and the non-STEM degree return by 55%. We then compare our model's performance against three alternative frameworks. First, a model with skill-neutral firm heterogeneity but no equity constraint can match log additive wages, but fails to generate worker-firm sorting or the hump-shaped size-pay relationship. Second, a model with productive complementarities between workers and firms can generate strong sorting patterns, but necessarily violates log additivity by introducing worker-firm interactions in wages. Third, a model with skill-varying labor supply elasticities can produce worker-firm sorting (while preserving log additivity), but cannot generate the non-monotonic relationship between firm size and wages in the data. Only our equity constraint framework can simultaneously accommodate all three empirical regularities, suggesting that pay compression within firms plays a fundamental role in shaping labor market equilibrium. Importantly, the equity constraint also has a strong basis in the theoretical and empirical literature highlighted above: we have not merely designed it to fit these empirical facts.

We also use our model to assess the welfare implications of internal pay equity constraints. Removing them would generate substantial benefits for high-skilled workers (increasing STEM graduate welfare by 42%) while harming low-skilled workers (reducing non-graduate welfare by 2%). These effects stem both from changes in wages and improved amenity matches (as low-skilled workers can now access the full set of firms[3]). This reflects a fundamental equity-efficiency trade-off: the removal of the equity constraint brings aggregate efficiency gains (through improvements in amenity matches), but exacerbates inequality. However, an alternative policy which prohibits selective hiring strategies (akin to mandating uniform hiring practices across firms) would bring both greater equity *and* efficiency gains. Non-graduate welfare would increase by 15%, while STEM graduate welfare would decrease by 5%; and at the same time, the elimination of workplace segregation would improve the

[3]In an alternative framework with job search frictions, this would manifest in reduced low-skilled unemployment.

quality of amenity matches.

We then use our framework to interpret market-level variation in firms' pay and recruitment strategies, over time, regions and sectors: these exercises shed new light on known empirical phenomena. First, at the aggregate level: given the growth in the relative supply and productivity of STEM workers (in the context of the Israeli tech boom), our model predicts greater adoption of selective hiring strategies. This should be reflected in greater pay dispersion across firms and heavier sorting of skilled workers to high-paying firms—and indeed, this is what the data show. This phenomenon may also be responsible for similar trends elsewhere, as documented by Card et al. (2013), Song et al. (2019) and Bonhomme et al. (2023).

Second, as a more demanding test, we exploit spatial labor market variation: according to the model, selective hiring strategies should be more pervasive in higher-skilled regions. Indeed, we show that regions with larger graduate employment shares (and larger skill expansions over time) exhibit greater firm pay dispersion and worker sorting (and greater increases in dispersion and sorting over time). These results speak to influential work by Dauth et al. (2022) and Card et al. (2025), who show that larger cities exhibit heavier sorting of workers to firms. Dauth et al. (2022) attribute this effect to scale economies in matching, but our model offers an alternative interpretation—arising from a quantity-quality trade-off in hiring.

Third, we study how pay equity constraints can shape firms' location choices. Recent evidence from Hazell et al. (2022) shows that multi-establishment firms face significant constraints in differentiating pay between workers in different locations; and we confirm this using Israeli data. However, in the face of these constraints, we argue that firms can use location choice as an additional tool to shape their skill mix—rather than just their wage policy. This generates a complementarity between pay and location strategies: high-paying firms concentrate in skilled regions while maintaining selective hiring, whereas firms with many locations choose moderate wages to accommodate a large and diverse workforce. These insights can be interpreted as a "skill analogue" to earlier work by Manning (2010), Hirsch et al. (2022) and Lindenlaub et al. (2024), who explore how firms trade off city size against wages in their location choices. Again, the data supports our predictions.

Finally, our framework can shed new light on differential wage returns within the public and private sectors. In many countries, the public sector offers lower returns to skill. This is typically attributed to tighter constraints on pay differentiation: see, e.g., Borjas (2002) on the US, and Mazar (2011) on Israel. However, we show that skill returns are no larger within *individual* private sector firms than in the public sector: i.e., equity constraints are similarly tight. Instead, what distinguishes the private sector is its fragmentation into many

independent firms (or "fissuring", in the language of Weil, 2014). This fragmentation facilitates larger returns to skill at the *aggregate* level, as firms adopt differential pay strategies, and high-skilled workers sort into high-paying firms—in line with our model. Conversely, the organizational unity of the public sector puts it at a disadvantage in the recruitment of skilled labor, especially in high-skilled cities (as in Propper and Van Reenen, 2010); and we confirm this empirically using spatial variation.

Our findings contribute to several strands of the labor literature. First, we add to a growing body of work documenting constraints on firms' ability to differentiate pay between their employees. Several papers show that firms cannot perfectly discriminate on workers' outside options: see e.g., Caldwell and Harmon (2019), Lachowska et al. (2022) and Di-Addario et al. (2023). Hazell et al. (2022) have explored constraints on pay discrimination by geography; and Amior and Manning (2020), Amior and Stuhler (2023) and Arellano-Bover and San (2023) study the implications of imperfect pay discrimination between natives and migrants.[4] Our focus here is pay compression among workers of different productivity.

We are not the first to explore the equilibrium implications of internal equity constraints: Romer (1984) and Akerlof and Yellen (1990) show how equity constraints can generate workplace segregation and unemployment of low-skilled workers. Beyond our empirical application using matched data, our key conceptual departure from these studies is to introduce wage-setting power, i.e., an imperfectly elastic supply of labor to the firm. This ensures that inclusive firms can maintain at least some high-skilled employment, despite offering them low pay. This is crucial to the profitability of the inclusive strategy, and hence the existence of a quantity-quality trade-off in equilibrium. In this respect, our model is more closely related to Manning (1994), who imposes a "company wage policy" (with firms constrained to paying a single wage to heterogeneous workers) on an equilibrium search model. Conceptually, our contribution is to partially relax this constraint, to allow for a limited degree of pay differentiation between workers within firms. This allows us to generate a log additive (AKM-type) wage structure, with distinct firm and worker effects. We then show how this framework can deliver a quantity-quality trade-off in hiring, which can help explain several empirical regularities in the literature. Finally, Frank (1984a,b) and Gola (2024) offer an alternative explanation for within-firm wage compression and workplace sorting, driven by workers' heterogeneous status concerns (rather than firms' wage-setting decisions); though this story does not deliver a quantity-quality trade-off.

Our work also builds on the extensive literature on sorting of high-quality workers to

[4]Our model is closely related to Amior and Stuhler (2023): they show how constrained pay differentiation between natives and migrants can generate workplace segregation and pay dispersion between ex-ante identical firms, whereas we apply this same idea to skill groups.

high-paying firms (Card et al., 2018; Dauth et al., 2022; Bonhomme et al., 2023; Haanwinckel, 2023). This sorting is often attributed to productive complementarities, but such complementarities are difficult to reconcile with the strong fit of log additive wage models. Instead, we provide a novel theoretical mechanism that can generate sorting in the absence of productive complementarities or even ex-ante firm heterogeneity. Our framework can also help explain why sorting patterns have intensified over time, and vary systematically across regions according to local workforce skill composition.

Finally, our paper speaks to the growing literature on the labor market consequences of domestic outsourcing (Abraham and Taylor, 1997; Weil, 2014; Goldschmidt and Schmieder, 2017; Deibler, 2022; Daruich et al., 2024; Gola, 2024). Interpreted through the lens of our model, outsourcing offers a means of escaping the pay constraint, by institutionally separating high and low-skilled employees. Many studies have focused on observable outsourcing events (to empirically identify causal effects), revealing wage losses for outsourced workers and (in the case of Deibler, 2022) gains for those who remain. But the phenomenon may be much broader, as firms may adopt selective or inclusive hiring strategies at the point of entry. Moreover, outsourcing is merely one manifestation of the quantity-quality trade-off: firms' rationing of low-skilled employees may also be absorbed through technological substitution in production, whether within defined roles (i.e., employing higher-quality workers to do given tasks) or through the adoption of alternative production processes.

In the next section, we present our theoretical framework and derive its key predictions. Section 3 describes our data, and Section 4 offers a quantitative assessment of our model: we document employment and wage patterns across the firm pay distribution, and calibrate the model to match these patterns. We then compare our model's performance against alternative frameworks, and assess key counterfactuals. In Section 5, we explore applications to temporal and spatial variation, as well as public-private sector differences; and we conclude in Section 6.

2 Equilibrium wage-setting model

We develop a simple equilibrium model of wage-setting, where firms are constrained in their ability to differentiate pay between workers of heterogeneous quality. As we will show, this pay equity constraint generates a novel trade-off between workforce quantity and quality, which can help shed new light on numerous labor market phenomena.

The economy consists of a continuum of firms (of measure k) and workers (measure n), who are either high or low-skilled.[5] In the baseline model, we assume firms are identical:

[5] As we show in Appendix D, the model is simple to extend to N skill types.

they produce a homogeneous output good, whose price is normalized to 1, with labor the sole factor of production. As in Card et al. (2018), firms choose skill-specific wages to maximize profit, and their wage-setting power arises from workers' idiosyncratic preferences over workplaces. We deviate from Card et al. (2018) by imposing a pay equity constraint: a strict within-firm limit on the wage differential between skill types.[6]

We begin by specifying labor supply. The utility of worker i of skill type $s = \{h, l\}$ in firm f takes the form:

$$u_{isf} = \varepsilon \log w_{sf} + a_{if} \tag{1}$$

where w_{sf} is the wage paid by firm f to type-s workers; and the a_{if} are idiosyncratic workplace amenity values, distributed type-1 extreme value. The supply of skill s labor to a firm offering wage w is then:

$$l_s(w) = \Omega_s w^{\varepsilon} \tag{2}$$

where ε is the elasticity of labor supply to individual firms (which is finite if firms have wage-setting power), and the intercept Ω_s depends on the aggregate skill s workforce, n_s, and competing wage offers:

$$\Omega_s = \left(\int_f w_{sf}^{\varepsilon} df \right)^{-1} n_s \tag{3}$$

We now turn to production. Like Card et al. (2018), we assume for simplicity that h- and l-type workers are perfect substitutes, but differ in efficiency units: h-types have (fixed) marginal product p_h, and l-types have marginal product p_l, where $p_h > p_l$. Firms choose wages w_s and employment l_s of each skill type $s = \{h, l\}$ to maximize profit π:

$$\max_{w_h, w_l, l_h, l_l} \pi(w_h, w_l, l_h, l_l) = (p_h - w_h) l_h + (p_l - w_l) l_l \tag{4}$$

subject to labor supply constraints:

$$l_h \leq l_h(w_h) \tag{5}$$

$$l_l \leq l_l(w_l) \tag{6}$$

and a pay equity constraint:

$$w_l \geq \phi w_h \tag{7}$$

The labor supply constraints (5) and (6) ensure that employment is bounded above by the labor supply curves: i.e., firms can only hire willing workers. The pay equity constraint

[6] Our hypothesis is not fundamentally tied to the monopsony model of Card et al. (2018). The key results can equally be derived from environments with job search frictions, such as Burdett and Mortensen (1998). We choose to build on Card et al. (2018) due to their model's analytical tractability.

(7) is our point of departure from standard monopsony models: firms cannot pay l-types less than a fraction ϕ of the h-type wage, where $\phi \leq 1$. Using the terminology of Weil (2014), this constraint may be interpreted in two ways: (i) as a "horizontal" equity constraint (with $\phi = 1$), where h and l-types are workers of different quality performing similar tasks (to different abilities), but firms cannot pay discriminate between them (a case explore by Manning, 1994); or (ii) as a "vertical" equity constraint (with $\phi < 1$), which limits the extent of pay differentiation between workers across the firm's hierarchy. The pay equity constraint can be microfounded using the efficiency wage model of Akerlof and Yellen (1990).[7]

The nature of labor market equilibrium depends on whether the equity constraint (7) binds or not. We will begin with the non-binding case, and then turn to the binding case.

2.1 Equilibrium if equity constraint does not bind

If the pay equity constraint does not bind, the labor supply constraints (5) and (6) must bind:

$$l_h^* = l_h(w_h) \tag{8}$$

$$l_l^* = l_l(w_l) \tag{9}$$

Intuitively, since firms set wages below marginal products, they will hire all workers who are willing to join them. For skill type $s \in \{h, l\}$, the optimal wage is then:

$$w_s^* = \frac{\varepsilon}{1+\varepsilon} p_s \tag{10}$$

which is a fixed mark-down on the marginal product p_s (determined by the labor supply elasticity ε). The wage differential will then equal the productivity differential:

$$\frac{w_l^*}{w_h^*} = \frac{p_l}{p_h} \tag{11}$$

From equation (11), the equity constraint will not bind if $\phi \leq \frac{p_l}{p_h}$.

[7] Following Akerlof and Yellen (1990), suppose an l-type worker's effort is given by $e_l = \min\left(\frac{w_l}{\tilde{w}_l}, 1\right)$, where $\tilde{w}_l = \phi w_h < p_l$ is a "fair wage" norm, and the corresponding productivity is $\tilde{p}_l = e_l p_l$. That is, workers only supply maximum effort ($e_l = 1$) if offered a wage exceeding the norm $\tilde{w}_l$. Under these assumptions, firms will never offer a wage below $\tilde{w}_l$. If they do so, profit per worker will equal $\tilde{p}_l - w_l = \left(\frac{p_l}{\tilde{w}_l} - 1\right) w_l$, which is *increasing* in the wage offer w_l; so firms can never benefit from offering a wage below the norm $\tilde{w}_l$. Intuitively, since firms set wages below marginal products, the savings on labor costs (from a wage cut below $\tilde{w}_l$) will not justify the associated productivity losses.

2.2 Equilibrium if equity constraint binds

Let β denote the bite of the pay equity constraint:

$$\beta \equiv \phi \frac{p_h}{p_l} \tag{12}$$

i.e., the ratio of the pay constraint ϕ to the productivity differential. The constraint binds if $\beta > 1$. Wages will then take log additive form:

$$\log w_{sf} = \eta_f + \lambda_s \tag{13}$$

The common firm effect η_f is chosen by firms, and is equal to $\log w_{hf}$ in the model. The skill effect $\lambda_s = I[s = l] \cdot \log \phi$ represents the fixed internal pay differential, which firms take as given.

In equilibrium, firms will adopt one of two pay strategies:

1. **Inclusive strategy (I).** Inclusive firms hire all willing workers, so the labor supply constraints bind for both skill types: i.e., $l_h^I = l_h\left(w_h^I\right)$ and $l_l^I = l_l\left(w_l^I\right)$. To accommodate both types, firms compress pay internally to satisfy the equity constraint, redistributing wages between h- and l-types (relative to the unconstrained optimum):

$$w_h^I = \frac{1 + \frac{1}{\beta} \cdot \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}}{1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}} \cdot w_h^* < w_h^* \tag{14}$$

$$w_l^I = \frac{\beta + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}}{1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}} \cdot w_l^* > w_l^* \tag{15}$$

   See Appendix B.1 for derivations.

2. **Selective strategy (S).** Selective firms hire all willing h-type workers, so the h-type labor supply constraint binds: i.e., $l_h^S = l_h\left(w_h^S\right)$. But they ration l-type employment. Since the l-type marginal product is fixed at p_l, rationing l-types only makes sense if the l-type wage w_l (which is fixed at ϕw_h if the equity constraint binds) exceeds p_l. If this is indeed the case, firms will optimally reject all l-type workers: i.e., $l_l = 0$.[8] And since selective firms hire only h-types, they will optimally offer them the unconstrained optimal wage: i.e., $w_h^S = w_h^*$. See Appendix B.2.

[8] More generally, if the marginal product p_l is decreasing in l_l (e.g., if skill types are imperfect substitutes or if there are diminishing returns to labor), optimal l-type employment for selective firms may strictly exceed zero, but still lie below the labor supply curve: i.e., $0 \leq l_l^S < l_l\left(\phi w_h^S\right)$. So even here, it remains true that selective firms ration l-type labor, and use it less intensively relative to h-type labor.

Though firms in this exposition are identical, they may choose different pay strategies in equilibrium. Let σ denote the equilibrium share of firms which choose the selective strategy. Equilibrium is uniquely determined, and can take one of two forms:

1. **Zero skill segregation.** The inclusive strategy yields strictly larger profit than the selective strategy: $\pi^I > \pi^S$. So all firms adopt the inclusive strategy, i.e., $\sigma = 0$. They pay the same wages, and hire equal shares of h- and l-type workers.

2. **Partial skill segregation.** Both strategies yield equal profit ($\pi^I = \pi^S$), so firms are indifferent between them. Since firms are identical, the adopted strategy of any given firm is undetermined; but the selective share σ is uniquely determined (there is a unique σ which equates π^I and π^S) and lies between 0 and 1. Selective firms pay high wages and recruit only h-types, and inclusive firms pay lower wages and recruit both h- and l-types; so skill types are partially segregated across firms.

   Note that firms' wage-setting power (i.e., a finite labor supply elasticity ϵ) is crucial to sustaining a partially segregated equilibrium, where inclusive firms offer lower pay. If labor supply were perfectly elastic, inclusive firms would not be able to maintain any h-type employment.[9]

As we show in Appendix B.3, the equilibrium σ can be expressed as:

$$\sigma = \begin{cases} 0 & \text{if } \beta < \frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}} - \alpha}{1-\alpha} \\ \tilde{\sigma}\left(\alpha, \beta, \varepsilon\right) & \text{if } \beta \geq \frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}} - \alpha}{1-\alpha} \end{cases} \tag{16}$$

where

$$\alpha \equiv \frac{p_h n_h}{p_h n_h + p_l n_l} \tag{17}$$

is the (exogenous) h-type aggregate output share, and the function $\tilde{\sigma}\left(\alpha, \beta, \varepsilon\right)$ solves the implicit equation:

$$\left(1 + \frac{1-\alpha}{\alpha - \tilde{\sigma}}\right)^{1+\varepsilon} = \left(1 + \beta\frac{1-\alpha}{\alpha - \tilde{\sigma}}\right)^{\varepsilon} \tag{18}$$

Equation (16) shows that the equilibrium selective share σ is fully determined by three parameters: the h-type output share α, the constraint bite β, and the labor supply elasticity ε. If the constraint bite is sufficiently weak, i.e., if $\beta < \frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}} - \alpha}{1-\alpha}$, the selective share σ is fixed at zero—and invariant to α, β and ε. But if β exceeds this threshold, σ is strictly increasing in all three parameters, in line with equation (18). The selective share σ is bounded above

[9] This explains why low-paying inclusive firms do not exist in the model of Akerlof and Yellen (1990): though they impose a similar internal equity constraint, they assume a competitive labor market.

by the h-type output share α: as σ converges to α, we move towards perfect skill segregation, with selective and inclusive firms exclusively employing h- and l-types respectively, and with all firms paying the unconstrained optimal wages (in equation (10)) and earning equal profit.

2.3 Comparative statics: Impact of equity constraint

If the equity constraint binds, firms face a trade-off between quantity and quality in hiring: by offering higher pay, selective firms can hire more h-type workers, but must ration l-type employment. This trade-off has important implications for pay dispersion, workplace segregation, earnings inequality and efficiency, which we now discuss:

Proposition 1. *An equity constraint with sufficient bite β generates:*

(a) Pay dispersion even among productively identical firms.

(b) Rationing of l-type workers by high-paying firms and hence workplace segregation.

(c) Compression of skill wage differentials, but no change in aggregate earnings.

(d) Reduction in expected amenity match quality, for both skill types. Since aggregate earnings and output are unchanged, this means the equity constraint is socially inefficient.

If the equity constraint has sufficient bite, and specifically if $\beta > \frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}} - \alpha}{1-\alpha}$, the equilibrium selective share σ will exceed zero. Selective firms will offer a high wage w_h^S, and inclusive firms a low wage w_h^I, to identical workers: this is part (a) of the proposition. This equilibrium is sustained by a quantity-quality trade-off: high-paying (selective) firms recruit more h-types, but this strategy compels them to ration l-type labor; and the equilibrium σ ensures that firms are indifferent between strategies.[10] Therefore, h-types disproportionately concentrate in selective firms (which offer them higher pay), and l-types only work for inclusive firms (as selective firms deny them employment): this is part (b).

Next, consider the implications for wage equity. In an equilibrium with zero skill segregation, it is clear that a binding equity constraint must compress wage differentials between skill types: firms simply respond by redistributing earnings between skill types, in line with equations (14) and (15). However, if β increases beyond the $\frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}} - \alpha}{1-\alpha}$ threshold, firms begin to adopt the selective strategy and ration l-type employment. Despite reduced *within*-firm wage differentials, growing segregation *between* firms gradually erodes the pay compression

[10] Equilibrium pay dispersion among identical firms is reminiscent of Burdett and Mortensen (1998). In both their model and ours, pay dispersion arises from a trade-off in the wage-setting decision, with different strategies yielding identical profit. For Burdett and Mortensen, this trade-off arises from the standard quantity motive of a non-discriminating monopsonist, in the context of on-the-job search: larger pay reduces profit per worker, but increases firm size. In our model, there is an additional quality motive in the trade-off, which arises from the binding pay constraint: firms use pay to shape their workforce composition, and not just workforce size. This quality motive delivers equilibrium pay dispersion even without on-the-job search.

effects. In the limit, as the bite β becomes very large, the labor market converges to perfect skill segregation (with $\sigma = \alpha$), with earnings differentials matching the unconstrained equilibrium. Up to this limit however, we show in Appendix B.4 that expected wage differentials are always narrower than in a counterfactual with no binding equity constraint. We also show that β has no effect on aggregate earnings, so these equity effects involve redistribution of earnings between workers alone. This is part (c) of the proposition.

We now turn to welfare. In this model, welfare depends not only on wages, but also on workplace amenities. If β exceeds the $\frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}}-\alpha}{1-\alpha}$ threshold (so the selective share σ exceeds zero), we show in Appendix B.5 that the equity constraint reduces the expected value of amenity matches, for both skill types. For l-types, this is because they are denied access to selective firms—and therefore have fewer firms to choose from.[11] The amenity loss may be so large that expected l-type utility decreases (despite the increase in earnings). For h-types, the amenity loss is a consequence of firm pay dispersion: h-types are willing to sacrifice amenity match quality to ensure employment at high-paying selective firms. Since aggregate earnings, profit and output are unchanged[12], these amenity losses imply that the equity constraint is socially inefficient: this is part (d). This result is in stark contrast to alternative models with firm-worker complementarities in production, where sorting of high-skilled workers to high-paying firms is associated with efficiency *gains*.

2.4 Implications for firm size

We next consider the implications for firm size:

Proposition 2. *An equity constraint with sufficient bite β generates:*

(a) A negative relationship between log firm size and pay, if firms are ex-ante identical.

(b) An initially positive and concave (and possibly hump-shaped) relationship, if there is skill-neutral heterogeneity in firm productivity.

We begin with part (a). In the baseline model with identical firms, if the equity constraint has sufficient bite (such that the selective share σ exceeds zero), selective firms will offer higher pay, but will have lower employment overall. This is a necessary consequence of the quantity-quality trade-off. Since firms are identical, the selective and inclusive strategies must deliver equal profit in equilibrium (the value of the selective share σ ensures this is the case). But selective firms employ more skilled workers, who individually generate larger

[11] In an alternative job search framework, this loss of access would manifest in higher unemployment for low-skilled workers, rather than lower-quality amenity matches: see Manning (1994).

[12] Output in this model is fixed by assumption, since workers are equally productive in all firms; so unchanged aggregate earnings implies unchanged profit.

profits. Therefore, to ensure equal profit, selective firms must employ fewer workers overall. See Appendix B.8 for a formal proof.

Of course, in practice, larger firms do typically pay higher wages; but the firm size premium is much smaller than what conventional monopsony models would predict, and the equity constraint can help explain why. To better match the data, we can extend the model to include skill-neutral heterogeneity in firm productivity. In a given firm f with firm-specific parameter x_f, suppose the h-type and l-type marginal products are equal to $p_{hf} = x_f p_h$ and $p_{lf} = x_f p_l$ respectively, where x_f is distributed log normally across firms.

In Appendix C, we show that log firm size is now initially positive and concave, and possibly hump-shaped, in log firm pay. Intuitively, skill-neutral heterogeneity in firm productivity introduces an *orthogonal* source of variation, which generates a countervailing *positive* correlation between firm size and pay. This positive correlation arises from the standard quantity motive: productive firms benefit more on the margin from larger employment, so they offer higher pay. However, the productivity parameter x_f makes no difference to the relative value of the selective and inclusive strategies (and hence, the equilibrium selective share σ is independent of x_f). Consequently, firm wage premia may now vary for two (orthogonal) reasons: (i) the choice of hiring strategy (selective firms offer higher pay) and (ii) through variation in productivity x_f (more productive firms offer higher pay).

Together, (i) and (ii) generate the concave relationship described by Proposition 1b. Since firms hire all willing h-type workers, the relationship between log h-type employment and log firm pay will simply trace the labor supply curve in (2): it will be positive and linear, with elasticity ε. However, the same is not true for l-type employment. Initially, for sufficiently low pay, the standard quantity motive dominates, and the slope will equal ε: higher-paying firms are more productive and recruit more workers. But higher up the pay distribution, the density of selective firms rapidly expands, and the quality motive plays a more important role: l-type employment is increasingly rationed, and this may even cause the firm size-pay relationship to turn negative (producing a hump-shaped relationship).

2.5 Implications for aggregate-level earnings inequality

Our model also delivers new insights on the determinants of aggregate-level earnings inequality. Increases in the relative productivity of h-types, i.e., $\frac{p_h}{p_l}$, and in their relative labor supply, $\frac{n_h}{n_l}$, make the selective hiring strategy more attractive; and this yields testable implications for workplace segregation and earnings differentials.

To guide our conceptual discussion—and the empirical analysis below—we will rely on a simple decomposition of skill wage differentials, derived from our model. Assuming the equity

constraint binds (i.e., $\beta > 1$), Appendix B.7 shows that the skill differential in expected log wages can be expressed as:

$$E[\log w_h] - E[\log w_l] = \underbrace{\log \frac{1}{\beta}\frac{p_h}{p_l}}_{\text{Within-firm}} + \underbrace{\frac{\sigma}{\alpha}\log\left(\frac{1-\sigma}{\alpha-\sigma}\right)^{\frac{1}{\varepsilon}}}_{\text{Between-firm}} \tag{19}$$

The first component on the right-hand side summarizes the contribution from within-firm pay differentials: from equation (12), notice that $\log \frac{1}{\beta}\frac{p_h}{p_l}$ is equal to $\log \frac{1}{\phi}$, where ϕ is the equity constraint. The second component summarizes the contribution from workplace segregation, i.e., the extent to which h-types are disproportionately employed by (high-paying) selective firms. Empirically, these components can be identified in two steps:

1. Estimate a log additive (AKM) model for wages, with worker and firm fixed effects.

2. Identify the first component using the mean differential in worker effects (between skill groups), and the second component by the mean differential in firm effects.

We now consider the determination of these components. Suppose for simplicity that the constraint bite β in equation (12) maintains its value in the face of any skill-biased changes in productivity (and the equity constraint ϕ adjusts to ensure this is the case: we leave a discussion of alternative assumptions to Appendix B.9[13]). Changes in relative h-type productivity $\frac{p_h}{p_l}$ will then be fully passed through to the within-firm component.

We next turn to the between-firm component. We have established above that workplace segregation can only exist in our model if the equity constraint binds. But the extent of segregation also depends on the relative h-type productivity $\frac{p_h}{p_l}$ and their relative labor supply $\frac{n_h}{n_l}$. In fact, holding the constraint bite β fixed, the impact of both can be summarized by a single parameter: the aggregate h-type output share α, as defined by (17). We make the following claim:

Proposition 3. *Assuming the equity constraint binds, and holding its bite β fixed, a larger h-type output share α increases (i) the equilibrium selective share σ and (ii) the between-firm component of the skill wage differential—as long as α is sufficiently large. Otherwise, α has no effect on the selective share and earnings inequality, and any productive benefits of larger α are shared equally between skill types.*

[13] If pass-through from relative productivity to within-firm pay differentials is only partial, this will amplify the positive impact of the h-type output share α on the selective share σ (and hence on the between-firm component). Intuitively, to the extent that firms cannot differentiate pay within firms, the quantity-quality trade-off becomes more acute; and workplace segregation increases in its stead.

See Appendix B.8 for a proof. If α is sufficiently small, such that $\frac{1-\alpha}{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}}-\alpha} < \beta$, it never makes sense for firms to adopt the selective strategy: there are not enough h-types (and/or they are not sufficiently productive) to justify rationing l-type employment. All firms will then offer the same wages to h- and l-types, defined by equations (14) and (15). Since the selective share σ is zero, there will be no workplace segregation and no between-firm component in the skill wage differential (19). In this region, the equity constraint compels firms to share any productive benefits of larger α equally between skill types.

But when the h-type output share α becomes sufficiently large, such that $\frac{1-\alpha}{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}}-\alpha} \geq \beta$, this sharing mechanism snaps: firms begin to adopt the selective strategy, and increasingly so as α grows. Since selective firms refuse to employ l-types, this expansion of σ ensures that only h-types capture the benefits from increases in α. This manifests through larger workplace segregation and a larger between-firm component in the skill wage differential.

The result that skill-biased productivity growth (in the presence of equity concerns) can generate workforce segregation is shared with the team formation framework of Gola (2024). But our model also delivers the same effect from increases in relative h-type labor supply. Both effects arise from changes in the quantity-quality trade-off facing wage-posting firms: a larger α increases the benefits of adopting the selective strategy (via the quantity and/or quality of h-type recruitment) and reduces the costs (associated with the rationing of l-type labor). The idea that a larger supply of skills can (perversely) increase earnings inequality is shared with Acemoglu (1998), but our story is very different—and centered around workforce segregation, rather than technical change.

3 Data and descriptive statistics

3.1 Data sources

Our analysis draws on Israeli administrative data covering the period 1990-2019. The core dataset, provided by the Central Bureau of Statistics (CBS), contains detailed employment records that link workers to firms. For each worker-firm match, we observe average monthly salary, industry classification, and an indicator for public sector employment. We restrict our main analysis to the private sector, where firms can plausibly adopt differential pay strategies (in line with our model). But in Section 5.4, we compare outcomes in the public sector—treating it as a "control" environment where this is not possible.

We link these records to comprehensive data on worker characteristics, including basic demographics and detailed information on education: we observe both highest degree completed and field of study. For our empirical analysis, we divide workers into three educa-

tion groups: (i) no college degree, (ii) non-STEM graduate, and (iii) STEM graduate. The STEM/non-STEM graduate distinction has become increasingly salient in recent decades (Altonji et al., 2016; Kirkebøen et al., 2016), and especially in the context of Israel and its tech boom. Though our model in Section 2 distinguishes between just two skill types, we show it is simple to extend to N types in Appendix D.

We also merge the employment records with detailed information on workplace locations. We group locations into 49 regions, based on Israel's "natural regions".[14] For some analysis, we rely on geographical identifiers from 20% samples of the Israeli census of 1995 and 2008; and for years from 2012, we rely on the Arnona (municipal tax) database for location information. We exploit spatial variation to test Proposition 3 (on the market-level determinants of firm pay dispersion and sorting), and to explore the implications for firms' location choices.[15]

3.2 AKM variance decomposition

In line with our model, we identify a firm's pay policy using the firm fixed effect in a log additive AKM wage specification, across workers i and years t:

$$\log w_{it} = \eta_{f(i,t)} + \lambda_i + \delta_t + \gamma X_{it} + \varepsilon_{it} \tag{20}$$

where $\eta_{f(i,t)}$ are firm effects (for the firm f employing worker i at time t), λ_i are worker effects, δ_t are year effects, and X_{it} includes time-varying controls.[16] The firm effects $\eta_{f(i,t)}$ are identified through worker mobility across firms.

Table 1 presents summary statistics and AKM variance decomposition results, both for our full sample and separately by education group, for years between 2010 and 2019. To address measurement error in the estimated firm effects, we implement a split-sample correction.[17] Panel A shows that the AKM model fits the data remarkably well, explaining 91.7% of the overall variance in log wages. The worker fixed effects account for the largest share of wage variance (61.5%), while firm effects contribute 8.3%, and the covariance between

[14] These have been defined by the Central Bureau of Statistics to ensure a high degree of uniformity in the demographic, economic, and social characteristics of the constituent population. We have merged the three smallest regions into neighboring regions, to ensure sufficient sample size for all empirical analysis.

[15] See Appendix H for further details on data definitions and processing.

[16] Following Card et al. (2018), we control for quadratic and cubic polynomials of age, centered around 40.

[17] This ensures our findings are not driven by statistical artifacts from estimation. Specifically, we randomly assign workers to two equally sized samples, "A" and "B", and estimate separate AKM models for each sample. For the firm effect variance in Table 1, we then use $\text{cov}(\eta^A_{f(i,t)}, \eta^B_{f(i,t)})$, where $\eta^A_{f(i,t)}$ and $\eta^B_{f(i,t)}$ are the firm effects estimated in the two samples. For the covariance between worker and firm effects, we use $\text{cov}(\lambda^A_i, \eta^B_{f(i,t)})$. Finally, for the worker effect variance, we first compute adjusted worker fixed effects using $\lambda^{adj}_i = \lambda^A_i + \eta^B_{f(i,t)} - \eta^A_{f(i,t)}$, and then use $\text{cov}(\lambda^A_i, \lambda^{adj}_i)$.

worker and firm effects explains 17.6%.[18] This indicates significant sorting of high-skilled workers to high-paying firms.

We also report results for an augmented specification with worker-firm interactions ("match effects"), which improves the R-squared by only 4 percentage points (95.7% versus 91.7% for the AKM model). As in Card et al. (2013), this small improvement suggests that a log-additive specification fits the data well, and match effects offer little additional explanatory power—consistent with our framework.

3.3 Returns to education

Panel B of Table 1 provides information on sample sizes and means. Our full sample consists of over 15 million worker-year observations between 2010 and 2019, with non-graduates constituting the largest group (59%), followed by non-STEM graduates (32%) and STEM graduates (9%). Average salaries are increasing in education: non-STEM graduates earn 0.28 log points more than non-graduates, and STEM graduates earn 0.74 log points more.

Following equation (19), we can decompose these raw wage differentials into within-firm and between-firm components, using our AKM estimates. Consider first the differential between non-STEM graduates and non-graduates: worker effects account for $\frac{0.11-(-0.11)}{9.25-8.97} = 79\%$ of this gap, while differences in firm effects account for $= \frac{0.03-(-0.03)}{9.25-8.97} = 21\%$.[19] Comparing STEM graduates and non-graduates, the between-firm component is even larger: it accounts for $30\% = \frac{0.19-(-0.03)}{9.71-8.97}$ of the raw differential. This illustrates the importance of worker-firm sorting in explaining the return to education. Interpreted through the lens of our model, the heavier sorting of STEM graduates is indicative of greater bite in their equity constraint (i.e., larger β): we will return to this point below.

4 Quantitative assessment of the model

In this section, we provide a quantitative assessment of our theoretical framework. We first document key empirical patterns in the Israeli labor market: the hump-shaped relationship between firm size and wage premia, heterogeneous employment patterns by education (and heavy worker-firm sorting), and log-additive wages. We then calibrate our model to match these patterns and compare its performance against alternative frameworks. We show that our equity constraint mechanism can simultaneously explain all three empirical regularities,

[18] These three components do not sum to the total R-squared, as our AKM model also controls for year effects and age polynomials.

[19] Looking at the table, the worker and firm effect differentials do not perfectly sum to the raw wage gap—but they are very close. This reflects the excellent fit of the AKM model.

while competing models fail to rationalize at least one key feature of the data. Our counterfactual analyses further illustrate the distributional implications of internal pay constraints and quantify their welfare effects across worker types.

4.1 Relationship between firm employment and pay premia

We begin in Figure 1 by plotting the relationship between log employment and firm AKM premia, i.e., η_f from equation (20), across firms. We group firms into 20 bins according to their AKM premia, with each bin containing an equal number of firms. The y-axis shows mean log firm employment in each bin, and the x-axis shows the mean firm premia, adjusted for measurement error using a split-sample correction.[20]

Panel A reveals a striking inverse-U shape. This pattern offers strong support for a quantity-quality trade-off in hiring, and is rationalized by Proposition 2. At low wage levels, employment increases with firm wage premia, consistent with the standard quantity motive: higher-paying firms are typically more productive and recruit more workers. However, the relationship is strongly concave—and even *decreasing* among the highest-paying firms. This is consistent with a rapidly expanding share of selective firms, which are prioritizing recruitment quality over quantity, and rationing lower-skilled employment.

We find similar patterns even within industries. In Panel B, we remove industry fixed effects from both the y-variable (log employment) and x-variable (firm premia); and the basic shape is preserved. This suggests it reflects fundamental trade-offs in firms' wage and hiring strategies, rather than simply industry-level differences in technology or skill requirements.

In Panels C and D, we exclude very small firms—with fewer than 5 employees. We continue to see a clear concave relationship in Panel C, but with no downward-sloping portion. However, the hump shape returns in Panel D when we remove industry effects.

The quantity-quality trade-off is more clearly revealed when we disaggregate employment by education. Figure 2 plots the relationship between firms' log education-specific employment (non-graduates, non-STEM graduates and STEM graduates) and their AKM premia. For non-graduates, we observe a strong hump-shaped pattern, with employment declining sharply at higher wage premia—and this time in all four panels. Employment of non-STEM graduates is strongly concave (but less so than for non-graduates), and has no clear downward-sloping portion. In contrast, STEM employment increases close to linearly with firm wage premia, exactly as predicted by our model: firms never ration high-skilled employment, so the green line simply traces out the isoelastic labor supply curve.

[20]To implement this correction, we randomly assign workers to two equally sized samples ("A" and "B"), and estimate the AKM model separately using each sample. We group firms into the 20 bins according to their sample A premia; and for each bin, we report the mean of the sample B premia on the x-axis.

Finally, at least for the within-industry estimates in Panel D, notice that the slopes (by education) are similar at the bottom of the firm pay distribution. This is consistent with the model, under the assumption of a common labor supply elasticity ε: see Section 2.4. Intuitively, among low-paying firms, the standard quantity motive dominates, and there is little skill rationing; so the employment slopes will be close to ε for all skill groups. We will rely on this insight below to quantify the model.

Hump-shaped employment is not particular to the Israeli context. In Appendix I, we replicate our analysis using the Veneto Worker History (VWH) dataset, which contains detailed employer-employee linked administrative records for Italy's Veneto region. As Figure A1 shows, we find a similar inverse-U relationship between firm size and firm wage premia. This suggests that the quantity-quality trade-off is a general phenomenon, arising from fundamental constraints on firms' wage-setting, rather than from country-specific institutions or policies.

Concavity in the firm size-pay premium relationship is not unique to our model. As Kline (2025) emphasizes, its shape will depend on distributional assumptions on workers' outside options (or the specification of utility). For simplicity, we have assumed an isoelastic labor supply elasticity ε in Section 2. But if, for example, workers' outside options are distributed according to a shifted power function (as in Card et al., 2018), the elasticity of labor supply will be decreasing in firm pay. But while this can account for concavity in the relationship, it cannot rationalize a hump shape; and it cannot explain why it is specifically lower-skilled workers who drive the concavity. Our model can explain both these features, as high-paying selective firms ration their employment of low-skilled labor in a quantity-quality trade-off.[21]

4.2 Log additivity of wages

If the equity constraint binds, firms in our model will share wage premia proportionally between skill types: i.e., wages will be log additive. We test this assumption by estimating the AKM model (20) separately by education group, and recovering group-specific premia. Figure 3 plots these group-specific premia against the aggregate (i.e., full sample) firm premia, across 20 bins (ordered by the aggregate premia). The bins are defined separately by education group, and contain equal numbers of group-specific workers; since STEM workers sort into higher-paying firms, the green bins are located more to the right.[22] Group-specific

[21] See Section 4.4 for further comparison between our model and alternatives from the literature.

[22] As before, we correct for measurement error using a split-sample method. We begin by randomly dividing workers into two samples: "A" and "B". For each sample, we estimate AKM firm premia using all workers ("aggregate premia") and separately by education group. For the non-graduate group (in blue), we split firms into 20 bins with equal numbers of non-graduate workers, according to their sample A aggregate premia. Along the x-axis, we report the mean sample B aggregate premia; and along the y-axis, we report

and aggregate premia are normalized to zero for firms with mean (employment-weighted) aggregate premia. If wages are log additive, the firm premia should then be identical across groups: i.e., the group-specific premia should increase one-for-one with the aggregate premia, and should line up perfectly on the 45-degree (dashed) line. Looking at Figure 3, the data are remarkably close to the dashed line, for all three education groups. Panel B shows the same patterns manifest within industries. These results are consistent with Card et al. (2018), who find that relative pay premia (of graduates to non-graduates) are very similar in high and low-value added firms.

There is a clear tension between Figure 3 and the sorting patterns in Figure 2, which has previously been highlighted by Bonhomme et al. (2019) and Kline (2025). If high-paying firms demand disproportionately more high-skilled workers, we might expect these firms to compensate them disproportionately—but Figure 3 shows otherwise. We argue that this tension can be resolved by an internal equity constraint, which delivers log additive wages and simultaneously compels high-paying firms to ration low-skilled employment. We elaborate on this point in Section 4.5 below.

4.3 Model quantification

The qualitative patterns above offer compelling support for our interpretation of the data. But we also fit the data quantitatively to our very parsimonious model. We study a specification with skill-neutral heterogeneity in firm productivity and three skill types, corresponding to non-graduates, non-STEM graduates and STEM graduates: we denote these l, m and h, respectively. The quantification exercise, detailed in Appendix E, identifies key parameters by matching observable moments in the data.

Table 2 summarizes the target moments and resulting parameter estimates. We identify the labor supply elasticity ($\varepsilon = 3.78$) using the relationship between log firm size and AKM firm effects at the bottom of the wage distribution: as we explain above, our model predicts that the standard quantity motive dominates in this region (so the firm size slope will approximate the true labor supply elasticity). And we calibrate the productivity variance ($\nu = 0.02$) to match the variance of AKM firm effects.

The equity constraint parameter $\phi_l = 0.59$ represents the (binding) wage ratio of non-graduates to STEM graduates; and $\phi_m = 0.79$ represents the wage ratio of non-STEM graduates to STEM graduates. These are identified by education differentials in mean AKM worker effects, from Table 1. Given our estimates of relative productivity (itself identified by the extent of sorting), these constraints indicate substantial compression of wage differentials

the mean sample B non-graduate premia. The red and green dots repeat this exercise for non-STEM and STEM graduates, respectively.

within firms. The equity constraint compresses the STEM degree return by 31% (relative to the STEM graduate v non-graduate productivity differential) within firms; and it compresses the non-STEM degree return by 55%.[23]

In this model with three skill types (a special case of the N-type model in Appendix D), firms pursue one of three strategies in equilibrium, which follow a hierarchical structure[24]: (i) a fully inclusive L-strategy, where firms hire all willing workers; (ii) an intermediate M-strategy, where firms hire only m- and h-types; and (iii) a highly selective H-strategy, where firms hire only h-types. As before, these strategies differ in the optimal h-type wage: we denote these as w_h^L, w_h^M and w_h^H respectively. Wages of the other skill types are then fixed by the equity constraints, ϕ_l and ϕ_m. In our calibration, $\sigma^H = 6\%$ of firms pursue the H-strategy (and hire only STEM graduates), $\sigma^M = 16\%$ adopt the M-strategy (and hire both STEM and non-STEM graduates); and the remaining $\sigma^L = 78\%$ adopt the fully inclusive L-strategy and employ all willing workers. Table 2 reports estimated h-type wage differentials between firms (of given productivity): M-strategy firms pay 30% less than H-strategy firms, and L-strategy firms pay 43% less.

4.4 Comparison with alternative models

To evaluate the performance of our framework (Model 1), we compare it to three alternative specifications: an equivalent model with skill-neutral firm heterogeneity but no equity constraint (Model 2); a model with productive complementarities between worker skill and firm quality (Model 3), in the spirit of Becker (1973); and a model with skill-varying labor supply elasticities (Model 4), as explored by Kline (2025). Figures 4-6 demonstrate that only our baseline model can match all the key empirical patterns documented in Figures 1–3.

Firm size-wage relationship. Figure 1 documents a striking inverse-U relationship between firm size and wage premia in the data. But Figure 4 shows that only our equity constraint model (Model 1) can successfully reproduce this pattern. Alternative models (2–4) all predict monotonically increasing relationships between firm size and wage premia, failing to generate the decreasing segment at higher wage levels. This hump shape emerges naturally in our model through the quantity-quality trade-off, which originates from the equity constraint.

[23] We compute compression of the STEM degree return as $\frac{\frac{1}{\phi_l}-1}{\frac{p_h}{p_l}-1} = 31\%$, and compression of the non-STEM degree return as $\frac{\frac{\phi_m}{\phi_l}-1}{\frac{p_m}{p_h}\cdot\frac{p_h}{p_l}-1} = 55\%$.

[24] This is the case if the equity constraint binds more for l-types than m-types: i.e., if $\beta_l > \beta_h$, where $\beta_s = \phi_s \frac{p_h}{p_s}$. We validate this assumption ex post.

Employment by education and worker-firm sorting. Figure 2 reveals that the hump shape is driven by low-skilled employment: in contrast, high-skilled employment increases monotonically with wages. These patterns imply sorting of high-skilled workers to high-paying firms (a well-known empirical regularity); but Figure 2 shows this sorting has a very particular character, with high-paying firms apparently rationing low-skilled workers. Figure 5 shows that only Model 1 can replicate these patterns. Though Models 3–4 do generate positive sorting, they both predict monotonically increasing relationships for all skill groups—and miss the sharp decline in low-skilled employment among high-paying firms.

Log additive wages. Figure 3 shows that firm wage premia are close to proportional across worker types, consistent with the log additive structure of AKM wage models. Our Model 1 delivers log additivity through the binding equity constraint, and Models 2 and 4 achieve the same result by assuming that firm heterogeneity is skill-neutral: see Figure 6. However, Model 3 violates log additivity, as the productive complementarities generate substantial worker-firm match effects.

These comparisons suggest that the equity constraint represents a fundamental feature of wage-setting, which can help explain multiple empirical regularities. While alternative models can match one or two of these results, only our framework can simultaneously account for all three. Aside from explaining these empirical regularities, the equity constraint has intuitive appeal: as we argue in the introduction, it has a strong basis in both the theoretical and empirical literature.

4.5 Resolution of empirical puzzles

The results above offer a resolution to two important empirical puzzles in the literature on firm wage-setting: on sorting of workers across firms and the firm size wage premium.

First, our framework can reconcile the heavy sorting of high-quality workers to high-paying firms with the remarkable empirical fit of log additive wage specifications. The tension between the two has previously been highlighted by Bonhomme et al. (2019) and Kline (2025). The most natural explanation for sorting is productive complementarities (as in Becker, 1973); but as Model 3 above shows, this structure fails to generate log additive wages. Our analysis here is very much in the spirit of Kline, who considers whether alternative models can resolve this apparent puzzle. Kline notes that a model with differential labor supply elasticities ε (like Model 4) can generate positive sorting, but argues that a larger high-skilled ε is difficult to rationalize. One might alternatively attribute positive sorting to workplace amenities: if more productive firms have better amenities, and if high-skilled

workers place greater value on these amenities, they may sort differentially into productive firms. But as Kline notes, we lack clear evidence for differential valuations of workplace amenities. Kline also offers another interpretation: wages may function as a screening device. If high-skilled workers have better outside options, and firms cannot condition wages on skill, higher wage offers may differentially attract high-skilled workers (see also Weiss, 1980). As Kline shows, this can yield comparable sorting patterns to Model 4. However, like Model 4, this model cannot reproduce the very particular character of sorting we identify in Figure 2, with a sharp decrease in low-skilled employment among high-paying firms. To account for this pattern, we require an explanation for why high-paying firms may choose to *ration* low-skilled employment: the equity constraint provides exactly this, through the quantity-quality trade-off.

Second, the model can help explain why firm size wage premia are significantly smaller than standard monopsony models would predict—typically showing only a 0.05 log wage increase per log point increase in firm size (Sokolova and Sorensen, 2021; Bloesch and Larsen, 2023). Conventional monopsony models require an implausibly elastic labor supply to individual firms to generate such small premia, with ε in the region of 20 (the inverse of 0.05). Our model, however, naturally produces a concave or even hump-shaped relationship between firm size and wages through the quantity-quality trade-off: selective firms offer higher pay, but ration their low-skilled employment; and this implies a much smaller wage return to firm size. A simulation of our model yields a firm size premium of 0.12 (from a regression of AKM firm premia on log employment): this is substantially below the value implied by conventional monopsony models ($1/\varepsilon = 0.26$, for our ε of 3.78), and much closer to our empirical estimate in Israeli data (0.074).[25]

It is notable that our estimate of the labor supply elasticity ($\varepsilon = 3.78$) aligns closely with recent estimates identified from within-firm variation. For example, tracing out the response to firm-level productivity shocks, Lamadon et al. (2022) and Kroft et al. (2020) find that employment grows 4-5 times as much as wages. This is consistent with our model: a skill-neutral productivity shock should not induce firms to adjust their hiring strategy (between selective or inclusive approaches), but only to adjust on the *quantity* margin.[26] Hence, the response to such a shock should be fully captured by the ε parameter. In contrast, when studying the *cross-sectional* distribution of firms (as in Figures 1 and 2), variation in hiring

[25] One can alternatively account for a small firm size premium by introducing a third factor which generates firm-level variation in employment independently of wages: in particular, Bloesch and Larsen (2023) propose a role for recruitment expenditures. However, our model makes a stronger prediction: that the relationship between firm size and pay is concave and potentially non-monotonic, and that this non-linearity is fully attributable to lower-skilled workers—just as we observe empirically.

[26] See the discussion in Section 2.4 and Appendix C.

strategy becomes much more salient—and employment may even be *decreasing* in wages among the highest-paying firms, as we show empirically.

4.6 Counterfactual analysis

Having established the empirical appeal of our model, we now quantify the welfare implications of internal equity constraints by conducting two counterfactual exercises. First, we examine the consequences of removing the equity constraint entirely, i.e., the reverse experiment of Proposition 1. Second, we consider a scenario where firms are prohibited from adopting the selective strategy—creating an environment similar to the public sector, which we explore empirically below. Table 3 presents changes in expected utility by education in these counterfactuals, decomposed into contributions from expected log wages and amenity match quality. Note we weight utility and amenity effects by $\frac{1}{\varepsilon}$ for this exercise, to ensure they are in log wage units: see equation (1). We leave technical derivations of these effects to Appendix F.

Panel A shows the effects of removing the equity constraint. Consistent with Proposition 1c, we see greater wage inequality between skill groups: firms now set wages independently for each group, and no longer redistribute rents between them. STEM graduates enjoy the largest wage gains (0.33 log points), with a smaller increase for non-STEM graduates (0.04), and significant losses for non-graduates (-0.09). At the same time, all three groups benefit from improved amenity match quality, consistent with Proposition 1d. Intuitively, in the counterfactual, high-skilled workers no longer need to sacrifice amenity match quality to ensure employment at high-paying selective firms; and low-skilled workers are no longer rationed by selective firms, so have more firms to choose from. However, the amenity gain for non-graduates (0.07 log points) is not sufficient to offset their wage losses; so their expected utility falls. These results highlight the equity-efficiency trade-off inherent in the equity constraint: its removal brings aggregate efficiency gains (amenity matches improve, with no change in aggregate output), but exacerbates inequality.

However, an alternative policy which prohibits selective hiring (and rationing low-skilled labor) can bring both greater equity *and* efficiency. We explore this counterfactual in Panel B. In equilibrium, conditional on their productivity, all firms offer the same wages (in line with the inclusive strategy) and redistribute rents between their high and low-skilled employees. The welfare impacts are therefore reversed: non-graduates enjoy large wage gains (0.08 log points), non-STEM graduates also benefit (0.02), but expected STEM wages contract by 14 log points. However, expected amenities still increase for all groups, just as in Panel A and for identical reasons: high-skilled workers benefit from reduced firm pay dispersion, and

low-skilled workers from access to all firms. Since aggregate output is unchanged (due to linear technology and full employment), we therefore have efficiency gains—alongside the improvement in equity.

The second counterfactual provides a useful theoretical benchmark for interpreting the public sector labor market. Given its organizational unity, the various administrative units of the public sector are unable to adopt differential pay strategies; so effectively, these units are compelled to adopt the inclusive strategy. As Panel B shows, this generates better outcomes for lower-skilled workers. But in practice, the public sector must compete with private firms; and an inclusive pay strategy makes it harder to attract high-skilled talent. We explore these questions empirically in Section 5.4 below.

5 Applications

5.1 Temporal variation in quantity-quality trade-off

Until now, we have focused on empirical variation across the distribution of *firms*. But the model also yields testable implications for *market-level* variation. Proposition 3 predicts that, in markets where high-skilled workers are more numerous and/or more productive (i.e., larger h-type output share α), more firms will adopt the selective strategy in equilibrium—and compromise on hiring quantity in favor of quality. As the labor market becomes more segregated, we should then expect greater dispersion in firm wage premia, with heavier sorting of high-skilled workers to high-paying firms.

The Israeli tech boom provides a natural setting to test these predictions. Table 4 shows large growth in the STEM workforce, in proportion terms: the STEM graduate employment share grew from 6.5% in the 1990s to 9.1% in the 2010s, with most of the change coming earlier in the sample. Over the same period, wage differentials between STEM graduates and non-graduates grew rapidly from 0.29 to 0.74. In contrast, the non-STEM graduate share has been close to flat, at around 30%; and the return to non-STEM degrees grew much more mildly, from 0.20 to 0.28.

In line with equation (19), we next study the contribution of firm effects to these changes in wage differentials—by estimating the AKM model separately for each decade. The table shows that most of the STEM return's growth can be attributed to increasing differentials in worker effects: i.e., within-firm pay differentials or ϕ in the model. But $28\% = \frac{0.222-0.095}{0.740-0.286}$ is driven by growing differentials in firm effects: i.e., heavier sorting of STEM workers to high-paying firms. Interpreted through the lens of our model, this heavier sorting indicates that the bite of the equity constraint (β in the model) must have grown: i.e., though internal

pay differentials ϕ increased, they did not keep pace with productivity differentials. As a result, more firms found it optimal to pursue the selective high-wage strategy, despite the necessary compromise on hiring quantity. This can explain the growing dispersion of firm wage premia: looking at the bottom row of Table 4, the variance of firm effects increased from 0.029 in the 1990s to 0.037 in the 2010s.

These patterns of growing dispersion in firm pay and heavier sorting are not unique to Israel. Similar trends have been documented in several advanced economies: see Card et al. (2013) on Germany, Song et al. (2019) on the US, and Bonhomme et al. (2023) on Sweden. Our model offers a simple interpretation of this phenomenon, driven by a quantity-quality trade-off which has become ever more attractive to firms. This can also help explain the expansion of domestic outsourcing (as in e.g., Goldschmidt and Schmieder, 2017; Gola, 2024). But as we argue above, outsourcing is just one potential manifestation of our story: the exclusion of low-skilled workers by selective firms may also reflect genuine technological substitution in production (i.e., employing higher-quality workers to do given tasks) or the adoption of alternative production processes.

5.2 Spatial variation in quantity-quality trade-off

Of course, macro-level variation can be difficult to interpret. As a more compelling test, we next exploit regional variation in workforce composition within Israel—both in the cross-section, and using regional changes over time. We rely on workplace location data from 20% samples of the Israel census of 1995 and 2008 (note that much of the expansion of STEM employment occurs between these years). We match these records with AKM firm wage premia estimated for corresponding intervals (1993-1997 and 2006-2010). Appendix Table A2 documents regional variation in skill shares and wages in 1995 and 2008: mean graduate share grew from 0.39 to 0.49, and its standard deviation from 0.045 to 0.066.

We estimate two specifications:

$$y_{rt} = \beta_0 + \beta_1 \text{GradShare}_{rt} + d_t + \varepsilon_{rt} \tag{21}$$

$$y_{rt} = \beta_0 + \beta_1 \text{GradShare}_{rt} + d_r + d_t + \varepsilon_{rt} \tag{22}$$

where y_{rt} represents some outcome in region r at time t, GradShare_{rt} is the local graduate employment share, and d_r and d_t are region and year fixed effects respectively. The first specification, which excludes region fixed effects, leverages cross-sectional variation to compare regions with different graduate shares. The second specification exploits local changes in graduate shares within regions over time. We do not claim to be isolating "causal" vari-

ation in the local graduate share. Instead, we are using the graduate share as a proxy for the h-type output share α, which is increasing in both the relative employment and productivity of high-skilled workers. Our model makes predictions on how α relates to the firm pay distribution and worker sorting across firms, and we seek to test these predictions empirically.

We present our main results in Panel A of Table 5. First, in columns 1-2, we show that a larger regional graduate share is associated with significantly higher mean firm wage premia. This is consistent with more firms adopting a selective high-pay strategy. The estimated coefficients are very similar between regions (0.321) and within them (0.378), suggesting that the relationship is very robust: a 10 percentage point increase in local graduate share is associated with a 3-4% increase in average firm premia.

As more firms adopt the selective strategy, we also expect larger dispersion in firm pay premia—and greater sorting of high-skilled workers to high-paying firms. These predictions are validated by the remaining columns. A 10 percentage point increase in local graduate share is associated with a 0.006 increase in the variance of firm wage premia (columns 3-4) and a 0.06-0.07 point increase in the correlation between worker and firm AKM effects. Again, the results are remarkably similar in the between-region and within-region specifications. In Appendix Table A3, we replace the graduate share with distinct regional STEM and non-STEM shares: the effects are mostly driven by the former, especially in the fixed effect specifications. This is consistent with the Israeli tech boom playing an important role.

These results build on the influential work of Dauth et al. (2022), who find significantly more assortative matching between workers and firms in larger cities in Germany (which contributes to the city-size wage premium). They attribute this effect to more efficient job matching, i.e., increasing returns in the local matching technology. In this paper, we propose a complementary agglomeration story: larger regions typically have larger skill shares, and this encourages more firms to adopt selective hiring strategies if equity constraints bind. At least in the Israeli context, the empirical evidence appears to support our interpretation. In Panel B of Table 5, we find a positive effect of log regional employment on sorting in column 5 (consistent with Dauth et al. 2022[27]); but Panel C shows that the regional graduate share captures most of the effect in a horse-race between the two variables. Still, it is worth emphasizing that our regions are significantly smaller than those used by Dauth et al. (2022), and this may influence the results.[28]

Our results also speak to Card et al. (2025), who show that spatial variation in the college

[27]They estimate a coefficient of 0.061 in Germany, which is larger than our 0.028 estimate.

[28]Our regions had a mean population of 150,000 in 2008, compared to 400,000 for German labor market areas; though Dauth et al. (2022) show their results are robust to using finer spatial variation across German counties, with a mean population of 250,000.

wage premium is partly driven by differential sorting of workers to industries. To make this connection more explicit, we study the implications for the graduate wage premium in Table 6, separately for STEM degrees (in Panel A) and non-STEM degrees (Panel B). We continue to use estimating equations (21) and (22), but with different outcomes. In columns 1-2, the dependent variable is the log wage differentials between (STEM or non-STEM) graduates and non-graduates in region r: these are increasing in the local graduate share, both for between-region and within-region variation. We next disaggregate these effects using the decomposition of equation (19). In the between-region specification, Columns 3 shows that most of the variation in the graduate premium is driven by unobserved worker heterogeneity (associated with the AKM worker effects), consistent with Card et al. (2025); but worker effects matter less in the within-region specification (column 4). More importantly for our paper, columns 5-6 show that differential sorting of graduates to high-paying firms plays an important role, and especially for STEM graduates—amounting to 32% or 42% of the overall effect in columns 1-2, depending on specification. Compared to Card et al. (2025), we explore variation within regions over time (and not just between regions); and we offer a new story for the sorting effects. In Appendix Table A4 (columns 4-5), we show the sorting effects are mostly driven by the regional STEM share (rather than the non-STEM share)—and especially in the fixed effect specifications. Again, this speaks to the important role of the Israeli tech boom.

5.3 Spatial equity constraints and firms' location choices

Until now, we have focused on pay equity constraints between skill groups within firms. But recent evidence from Hazell et al. (2022) suggests that multi-establishment firms face analogous constraints on wage-setting across *regions* ("national wage-setting"). In this section, we offer evidence for this phenomenon in Israeli data: this validates the AKM framework's implicit assumptions on the spatial dimension. But it also yields new insights on how firms manage the quantity-quality trade-off. In our baseline model, firms shape their skill mix using their wage policies alone. But faced with a spatial equity constraint, firms' *location choices* also play an important role—and we demonstrate this empirically.

To analyze spatial wage-setting patterns, we first allocate workers to firm-region pairs. Using detailed address information from the Arnona (municipality tax) database, we observe the complete set of establishments of each firm. We aggregate these to the regional level, to define up to 58 potential units per firm—which we call "divisions".[29] While only 7% of firms

[29]We know the firm which employs any given worker, but we do not know which division (if the firm has multiple divisions). Our approach is to allocate workers to the closest division of their employer, using residence and workplace information from the Arnona database. See also Vilhuber (2018) and Card et al.

operate in multiple regions, these multi-region firms employ 53% of workers in our sample.

We begin by estimating divisional wage premia $\eta_{fr(i,t)}$, for individuals of fixed quality, by replacing the firm fixed effects with interacted firm × region effects in the AKM estimating equation (20):

$$\log w_{it} = \eta_{fr(i,t)} + \lambda_i + \delta_t + \gamma X_{it} + \varepsilon_{it} \tag{23}$$

We then study how divisional wage premia vary spatially within multi-region firms—and specifically, how they are influenced by regional market-level wage differentials. We estimate the following equation:

$$\eta_{fr} = \beta \bar{\eta}_r + \lambda_f + \varepsilon_{fr} \tag{24}$$

where $\bar{\eta}_r$ is the region-level mean (across all firms) of the divisional premia η_{fr}. Clearly, regressing η_{fr} on $\bar{\eta}_r$ would deliver a coefficient of 1; but by controlling for firm fixed effects λ_f, we can isolate the wage variation *within* firms. If $\beta = 1$ even conditional on firm effects, this would indicate that firms spatially differentiate pay one-for-one with regional market differentials (conditional on workers' quality). At the other extreme, a β of 0 would indicate that firms do not differentiate pay at all. This exercise builds heavily on Hazell et al. (2022), but we differ in using matched administrative data: this allows us to identify wage premia for otherwise identical workers, by controlling for worker fixed effects (and exploiting job switchers).

We present our estimates in Table 7. Our basic OLS estimate of β in column 1 is 0.18: i.e., in regions where market pay (conditional on worker quality) is 1 point higher, firms on average compensate their employees just 0.18 points more (again, conditional on worker quality). This is consistent with many firms implementing national pay policies, as in Hazell et al. (2022). But since $\eta_{fr(i,t)}$ and $\bar{\eta}_r$ are both generated variables, measurement error can cause systematic bias in OLS. In column 2, we use the regional graduate share as an instrument: as column 7 shows (and see also Table 5), it strongly predicts regional pay premia. The coefficient is now somewhat smaller, at 0.15.

These patterns have important implications for recruitment quality. If multi-region firms cannot spatially differentiate pay, they will tend to offer relatively low wages in high-wage (high-skilled) regions. Following the logic of our model, this means they will adopt more inclusive hiring strategies—and recruit lower-quality workers—in these regions than their competitors. To test this hypothesis, we estimate the following equation:

$$\text{GradShare}_{fr} = \beta \text{GradShare}_r + \lambda_f + \varepsilon_{fr} \tag{25}$$

(2025), who use distance to allocate workers to workplaces.

where GradShare_{fr} is the graduate employment share of the firm's division, and GradShare_{r} is the regional graduate share. Again, with no firm effects, the coefficient would necessarily equal 1. But controlling for λ_f, column 6 of Table 7 reports a β of 0.22. That is, multi-divisional firms typically recruit higher-skilled workers from higher-skilled regions. But since the coefficient is far below 1, they also recruit many fewer than their local competitors—consistent with our hypothesis.

This has important implications for firms' location choices. Consider a world with many regions, characterized by different skill compositions. Firms are free to open divisions in any region they choose, but must pay a fixed cost for each division—which is increasing in the number of divisions (e.g., due to managerial diseconomies). At the same time, suppose firms are subject to a spatial equity constraint, and must pay the same wage premia everywhere. The quantity-quality trade-off will then apply to quantity of regions, analogously to quantity of workers in the baseline model. Intuitively, firms which adopt selective hiring strategies will benefit disproportionately from locating in high-skilled regions (with larger α), where the benefits of the strategy are larger (and the costs smaller): a corollary of Proposition 2. But since these firms must offer high pay everywhere, they cannot profitably operate in low-skilled regions—so they open relatively few divisions: a corollary of Proposition 3. At the same time, if there is (skill-neutral) firm heterogeneity in productivity, the quantity motive will guide more productive firms to open more divisions—as they stand to benefit disproportionately more from higher output.

Figure 7 offers evidence for this hypothesis. First, Panel A plots the relationship between a firm's number of regions and its AKM premium, using the same bin structure (with split-sample adjustment) as before. This reveals a clear hump-shaped relationship, analogous to the employment patterns in Figure 1. At the same time, Panel B reveals that higher-paying firms systematically select into regions with larger graduate shares. These results are robust to excluding single-branch firms and controlling for industry effects.

These location choices can be interpreted as a form of "directed search", as firms—which face binding pay equity constraints—seek additional ways to manage the quantity-quality trade-off beyond wage policy alone. This is effectively a "skill analogue" to earlier work by Manning (2010), Hirsch et al. (2022) and Lindenlaub et al. (2024), who explore how firms trade off city size against wages in their location choices: given the conventional quantity motive, more productive firms disproportionately operate in larger cities—where they have access to more labor. But we show how local workforce composition (and not just city size) can be the driving force: the equity constraint generates a quantity-quality trade-off, which pushes high-paying selective firms towards high-skilled regions, even in the absence of productive complementarities.

5.4 Public sector wage returns

In many countries, the public sector offers lower returns to skill than the private sector. This is typically attributed to tighter constraints on pay differentiation (e.g., Borjas, 2002; Mazar, 2011). But our framework offers an alternative interpretation. We argue that individual private sector firms are no better at differentiating pay than the public sector: rather, the key distinction lies in the private sector's fragmentation into many independent firms. This fragmentation facilitates larger returns to skill at the *aggregate* level, as firms adopt differential pay strategies, and high-skilled workers sort into high-paying firms. That is, the public sector is an empirical analogue of the counterfactual with no selective strategy in Table 3.

To test this interpretation, we estimate the AKM model of equation (20) on the full sample, including both private and public sector employment. In our data, "firms" in the public sector identify different administrative units (with different tax codes). We save both the estimated firm effects (η_f) and worker effects (λ_i), corresponding to each individual worker. In Table 8, we then decompose the wage returns to education in each sector, in line with equation (19).

In Israel, the return to non-STEM degrees is slightly larger in the public sector—a consequence mostly of health and education professionals. But the return to STEM degrees is much larger in the private sector: 0.739 versus 0.504. Table 8 shows this difference is almost entirely driven by the *between-firm* component: within-firm wage differentials (i.e., in the worker effects) are remarkably similar across sectors, which suggests that constraints on internal pay differentiation are similarly. Instead, consistent with our hypothesis, the large private sector returns to STEM arise from the sorting of high-skilled workers into high-paying firms—a mechanism that is absent in the public sector, where organizational units cannot easily compete on pay.

Additional statistics in Table 8 provide further context. The public sector employs workers with higher average worker effects (0.152 versus -0.047 in the private sector), suggesting positive selection into public employment. Interestingly, the public and private sectors offer similar average firm effects (-0.023 versus 0.014), but the variance in the private sector is much larger (0.037 versus 0.018). Again, this reflects the fragmentation of the private sector—and the ability of firms to adopt distinct pay and hiring strategies in equilibrium.

Above, we have focused on the implications of the public sector's organizational unity in the face of equity constraints by *skill*. But analogous arguments apply to *spatial* pay constraints. Intuitively, the public sector functions as a large multi-region firm in Section 5.3. Since the firm cannot spatially differentiate its pay premia, it will adopt more inclusive hiring strategies (and recruit lower-quality workers) in high-skilled (high-wage) regions than

its local competitors—and the same is true of the public sector. This insight speaks to influential work by Propper and Van Reenen (2010), who show that English hospitals achieve worse health outcomes in regions with higher private sector wages (i.e., where high-quality health workers have more attractive outside options). They attribute this to pay regulation in the English health sector, which prevents spatial differentiation. But the message of our paper is that this is a much broader phenomenon, which affects unregulated sectors also. Again, what distinguishes the public sector is not its inability to differentiate pay, but rather its organizational unity.

We now test this idea empirically, by studying spatial variation in public sector pay and recruitment quality—relative to the private sector. Compared to Propper and Van Reenen (2010), we study the public sector as a whole (and not just hospitals) and focus on recruitment quality (rather than output). Building on the regional analysis in Section 5.2, we estimate the following two specifications:

$$y_{srt} = \beta_0 + \beta_1 \text{GradShare}_{rt} + \beta_2 \text{GradShare}_{srt} \cdot \text{Public}_s + d_{st} + \varepsilon_{srt} \tag{26}$$

$$y_{srt} = \beta_0 + \beta_1 \text{GradShare}_{rt} + \beta_2 \text{GradShare}_{srt} \cdot \text{Public}_s + d_{sr} + d_{st} + \varepsilon_{srt} \tag{27}$$

where y_{rt} represents some outcome y in sector s (private or public), region r at time t, GradShare_{rt} is the local graduate employment share, Public_s is a dummy taking 1 for the public sector, d_{sr} are interacted sector-region fixed effects, and d_{st} are interacted sector-time fixed effects. These specifications are identical to (21) and (22) above, except we now have double the observations (due to the sectoral disaggregation), and test for differential effects of the local graduate share on public sector outcomes using the β_2 coefficient. Equation (26) exploits variation *between* regions in graduate share, whereas (27) relies on changes in graduate share *within* regions.

We present our estimates in Table 9. The β_1 coefficient in columns 1-2 confirms that private sector firms offer larger wage premia in higher-skilled regions, consistent with Table 6. However, the β_2 coefficient shows this effect is fully offset in the public sector: i.e., public sector pay is not increasing in graduate share. Columns 3-4 show that dispersion in private sector firm premia is increasing in graduate share (as in Table 6); but again, this effect is significantly muted in the public sector. We see similar patterns for worker-firm sorting in columns 5-6: sorting is strongly increasing in graduate share in the private sector, but not in the public sector (at least in the cross-sectional specification in column 5).

To summarize, the estimates in columns 1-6 offer strong support for our hypothesis. Given its organizational unity, public sector units cannot implement the selective high-wage strategy that private firms use to attract high-skilled workers in high-skilled regions. Con-

sequently, the public sector exhibits systematically less dispersion in firm wage premia and weaker worker-firm sorting——and especially in high-skilled regions, where the private sector shows the strongest effects.

The remaining columns of Table 9 explore the consequences for recruitment quality across sectors. Given their heavy adoption of the selective strategy in high-skilled regions, we expect private sector firms to disproportionately capture high-skilled workers in these locations; and the estimates support this claim. In columns 7-8, the mean worker AKM effect is strongly increasing in regional graduate share in the private sector, but less so in the public sector (though the standard errors are large, and the β_2 interaction effect is not statistically significant). But when studying effects on non-STEM and STEM graduate employment shares (in columns 9-12), we see much stronger effects: public sector recruitment quality is close to flat in regional graduate share; whereas the private sector captures almost all of this regional skill variation.

6 Conclusion

It has long been argued that firms face significant constraints in their ability to differentiate pay by worker productivity, a claim supported by recent empirical work. In this paper, we show how these internal equity constraints generate a quantity-quality trade-off in hiring. Firms must choose between a selective hiring strategy—paying high wages to attract high-quality workers, while rationing lower-skilled employment—or an inclusive strategy—maintaining lower wages to employ a larger, more diverse workforce. Unlike in a conventional monopsony model, firms use higher pay to improve hiring quality, even at the cost of lower quantity.

This insight can help resolve several empirical puzzles in the labor literature. It provides a novel interpretation of the (empirically successful) log additive AKM wage model, and shows how log additivity can be reconciled with sorting of high-skilled workers to high-paying firms. It can also help explain why firm size premia are surprisingly small, without requiring implausibly elastic labor supply. And it provides a new perspective on why seemingly similar firms adopt different pay and hiring strategies, and why these differences vary systematically across regions and sectors.

Using detailed administrative data from Israel, we find strong empirical support for our model's predictions. We show that the relationship between firm size and wages follows an inverse-U shape (both on aggregate and within industries), whose concavity can be attributed to lower-skilled workers—consistent with high-paying firms rationing low-skilled employment. We show how our very parsimonious model can successfully match these em-

pirical patterns, in contrast to alternative monopsony models. We then use this model to explore key counterfactuals: eliminating the equity constraint would improve equity, at the expense of efficiency; but an alternative policy which prohibits the selective hiring strategy would benefit both equity and efficiency. Finally, we show how our model can shed new light on several other empirical labor market phenomena: on temporal and spatial variation in firm pay dispersion and firm-worker sorting, firms' location choices (in the context of spatial equity constraints), and wage returns and recruitment outcomes in the public sector.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica 67*(2), 251–333.

Abraham, K. G. and S. K. Taylor (1997). Firms' Use of Outside Contractors: Theory and Evidence. *Journal of Labor Economics 14*(3), 394–424.

Acemoglu, D. (1998). Why do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *Quarterly Journal of Economics 113*(4), 1055–1089.

Akerlof, G. A. and J. L. Yellen (1990). The Fair Wage-Effort Hypothesis and Unemployment. *Quarterly Journal of Economics 105*(2), 255–283.

Altonji, J. G., P. Arcidiacono, and A. Maurel (2016). The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects. In *Handbook of the Economics of Education*, Volume 5, pp. 305–396. Elsevier.

Amior, M. and A. Manning (2020). Monopsony and the Wage Effects of Migration. CEP Discussion Paper No. 1690.

Amior, M. and J. Stuhler (2023). Immigration, Monopsony and the Distribution of Firm Pay. CEP Discussion Paper No. 1690.

Arellano-Bover, J. and S. San (2023). The Role of Firms and Job Mobility in the Assimilation of Immigrants: Former Soviet Union Jews in Israel 1990-2019. IZA Discussion Paper No. 16389.

Becker, G. S. (1973). A Theory of Marriage: Part I. *Journal of Political Economy 81*(4), 813–846.

Bewley, T. F. (1999). *Why Wages Don't Fall During a Recession*. Harvard University Press.

Bloesch, J. and B. Larsen (2023). When do Firms Profit from Wage Setting Power? Unpublished.

Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023).

How Much Should We Trust Estimates of Firm Effects and Worker Sorting? *Journal of Labor Economics 41*(2), 291–322.

Bonhomme, S., T. Lamadon, and E. Manresa (2019). A Distributional Framework for Matched Employer Employee Data. *Econometrica 87*(3), 699–739.

Borjas, G. J. (2002). The Wage Structure and the Sorting of Workers into the Public Sector. NBER Working Paper No. 9313.

Breza, E., S. Kaur, and Y. Shamdasani (2018). The Morale Effects of Pay Inequality. *Quarterly Journal of Economics 133*(2), 611–663.

Burdett, K. and D. T. Mortensen (1998). Wage Differentials, Employer Size, and Unemployment. *International Economic Review*, 257–273.

Caldwell, S. and N. Harmon (2019). Outside Options, Bargaining, and Wages: Evidence from Coworker Networks. Unpublished.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics 36*(S1), S13–S70.

Card, D., J. Heining, and P. Kline (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics 128*(3), 967–1015.

Card, D., A. Mas, E. Moretti, and E. Saez (2012). Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. *American Economic Review 102*(6), 2981–3003.

Card, D., J. Rothstein, and M. Yi (2025). Location, Location, Location. *American Economic Journal: Applied Economics 17*(1), 297–336.

Daruich, D., M. Kuntze, P. Plotkin, and R. Saggio (2024). The Consequences of Domestic Outsourcing on Workers: New Evidence from Italian Administrative Data. Unpublished.

Dauth, W., S. Findeisen, E. Moretti, and J. Suedekum (2022). Matching in Cities. *Journal of the European Economic Association 20*(4), 1478–1521.

Deibler, D. (2022). The Effect of Outsourcing on Remaining Workers, Rent Distribution, and Inequality. Unpublished.

Di-Addario, S., P. Kline, R. Saggio, and M. Sølvsten (2023). It Ain't Where You're From, It's Where You're At: Hiring Origins, Firm Heterogeneity, and Wages. *Journal of Econometrics 233*(2), 340–374.

Dube, A., L. Giuliano, and J. Leonard (2019). Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. *American Economic Review 109*(2), 620–663.

Frank, R. H. (1984a). Are Workers Paid Their Marginal Products? *American Economic Review 74*(4), 549–571.

Frank, R. H. (1984b). Interdependent Preferences and the Competitive Wage Structure. *RAND Journal of Economics*, 510–520.

Giupponi, G. and S. Machin (2022). Company Wage Policy in a Low-Wage Labor Market. CEP Discussion Paper No. 1869.

Gola, P. (2024). The Pond Dilemma with Heterogeneous Relative Concerns. Unpublished.

Goldschmidt, D. and J. F. Schmieder (2017). The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure. *Quarterly Journal of Economics 132*(3), 1165–1217.

Haanwinckel, D. (2023). Supply, Demand, Institutions, and Firms: A Theory of Labor Market Sorting and the Wage Distribution. NBER Working Paper No. 31318.

Hazell, J., C. Patterson, H. Sarsons, and B. Taska (2022). National wage setting. NBER Working Paper No. 30623.

Hirsch, B., E. J. Jahn, A. Manning, and M. Oberfichtner (2022). The Urban Wage Premium in Imperfect Labor Markets. *Journal of Human Resources 57*(S), S111–S136.

Kirkebøen, L. J., E. Leuven, and M. Mogstad (2016). Field of Study, Earnings, and Self-Selection. *Quarterly Journal of Economics 131*(3), 1057–1111.

Kline, P. (2024). Firm Wage Effects. In *Handbook of Labor Economics*, Volume 5, pp. 115–181. Elsevier.

Kline, P. (2025). Labor Market Monopsony: Fundamentals and Frontiers. NBER Working Paper No. 33467.

Kroft, K., Y. Luo, M. Mogstad, and B. Setzler (2020). Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry. NBER Working Paper No. 27325.

Lachowska, M., A. Mas, R. Saggio, and S. A. Woodbury (2022). Wage Posting or Wage Bargaining? A Test Using Dual Jobholders. *Journal of Labor Economics 40*(S1), S469–S493.

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review 112*(1), 169–212.

Lindenlaub, I., R. Oh, and M. Peters (2024). Spatial Firm Sorting and Local Monopsony Power. NBER Working Paper No. 30637.

Manning, A. (1994). Labour Markets with Company Wage Policies. CEP Discussion Paper No. 214.

Manning, A. (2006). A Generalised Model of Monopsony. *Economic Journal 116*(508), 84–100.

Manning, A. (2010). The Plant Size-Place Effect: Agglomeration and Monopsony in Labour Markets. *Journal of Economic Geography 10*(5), 717–744.

Mazar, Y. (2011). Self-Selection of Employees Moving Between the Public and Private Sectors. *Israel Economic Review 9*(1).

Propper, C. and J. Van Reenen (2010). Can Pay Regulation Kill? Panel Data Evidence on the Effect of Labor Markets on Hospital Performance. *Journal of Political Economy 118*(2), 222–273.

Romer, D. (1984). The Theory of Social Custom: A Modification and Some Extensions. *Quarterly Journal of Economics 99*(4), 717–727.

Sokolova, A. and T. Sorensen (2021). Monopsony in Labor Markets: A Meta-Analysis. *ILR Review 74*(1), 27–55.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter (2019). Firming Up Inequality. *Quarterly Journal of Economics 134*(1), 1–50.

Vilhuber, L. (2018). LEHD Infrastructure S2014 files in the FSRDC. Technical report.

Weil, D. (2014). *The Fissured Workplace*. Harvard University Press.

Weiss, A. (1980). Job Queues and Layoffs in Labor Markets with Flexible Wages. *Journal of Political Economy 88*(3), 526–538.

Tables and figures

Table 1: Descriptive statistics and AKM decomposition by worker type

	All	Non-grads	Non-STEM grads	STEM grads
Panel A: AKM variance decomposition				
Var. log salary	0.42	0.32	0.43	0.47
AKM model (share):				
Var. worker effect	61.5	65.8	63.1	51.3
Var. firm effect	8.3	8.3	7.7	10.1
2 × Cov(worker, firm)	17.6	15.1	13.4	16.5
R-squared	91.7	90.0	91.1	91.1
Comparison match model (share:)				
R-squared	95.7	94.8	95.4	95.6
Panel B: Sample means and size				
Worker-years				
N.	15,306,750	9,002,405	4,908,020	1,396,325
Share N.	1.00	0.59	0.32	0.09
Av. log salary	9.12	8.97	9.25	9.71
Av. worker effect	0.00	-0.11	0.11	0.39
Av. firm effect	0.01	-0.03	0.03	0.19
Workers				
N.	2,867,339	1,677,599	954,309	235,431
Share N.	1.00	0.59	0.33	0.08
Firms				
N.	184,495			
Av. firm size	24.4			

Notes: Panel A presents variance decomposition results from an AKM model (one model for all worker types). We correct for measurement error using a split-sample procedure (see text). The final row reports the R-squared of a model with interacted firm-worker fixed effects. Panel B presents the number of observations and averages of relevant variables for worker-years, workers, and firms. Sample consists of private sector firms between 2010 and 2019.

Table 2: Quantification of model parameters

Moments			Parameters		
Moment	Value	Source	Parameter	Value	Equation
ε_{bottom}	3.61	Elasticity of firm size, for bottom 25% firm pay	ε	3.78	-
V_{AKMf}	0.035	Variance of firm effect	ν	0.02	-
ϕ_m	0.79	m-type person effect	σ^M	0.16	(D6)
ϕ_l	0.59	l-type person effect	σ^L	0.78	(D5)
$E[\log w_m]$ $-E[\log w_h]$	-0.40	Expected m-type log wage, relative to h-type	$\frac{w_h^M}{w_h^H}$	0.70	(D2)
$E[\log w_l]$ $-E[\log w_h]$	-0.75	Expected l-type log wage, relative to h-type	$\frac{w_h^L}{w_h^H}$	0.57	(D1)
$\frac{n_m}{n_h}$	3.81	Relative m-type employment	$\frac{\phi_m^\varepsilon \Omega_m}{\Omega_h}$	5.73	(D4)
$\frac{n_l}{n_h}$	6.85	Relative l-type employment	$\frac{\phi_l^\varepsilon \Omega_l}{\Omega_h}$	14.69	(D3)
			Implied parameters		
			$\frac{p_m}{p_h}$	0.50	(C22)
			$\frac{p_l}{p_h}$	0.31	(C22)
			β_m	1.58	(C36)
			β_l	1.91	(C36)

Notes: This table shows the empirical moments used for model calibration (left columns) and the resulting parameter estimates (right columns). The labor supply elasticity (ε) is identified from the relationship between firm size and wage premia at the bottom of the firm pay distribution. The productivity variance (ν) is calibrated to match the variance of firm effects. Strategy-specific wages and supply intercepts are recovered from average wage differentials and employment ratios. σ^M and σ^L are respectively the shares of firms adopting the M-strategy (i.e. hiring only m- and h-type workers) and L-strategy (hiring all skill types). See Appendix E for more details.

Table 3: Welfare effects of counterfactuals

Type	Exp log wage	Exp amenity	Exp utility
Panel A: No pay-equity constraint			
Non-graduates	-0.091	0.067	-0.023
Non-STEM graduates	0.039	0.033	0.072
STEM graduates	0.331	0.091	0.422
Panel B: No selective strategy			
Non-graduates	0.083	0.067	0.151
Non-STEM graduates	0.024	0.033	0.057
STEM graduates	-0.143	0.091	-0.051

Notes: This table presents welfare changes from two counterfactual scenarios. Panel A shows what happens if we eliminate the pay equity constraint, allowing firms to set wages of skill types independently. Panel B shows what happens if we prohibit the selective pay strategy, requiring all firms to employ workers of all skill types. Worker types are defined by education: STEM graduates (type-h in the model), non-STEM graduates (type-m), and non-graduates (type-l). Changes in expected utility are decomposed into changes in expected log wages and expected amenity matches. Note we weight utility and amenity effects by $\frac{1}{\varepsilon}$ for this exercise, to ensure they are in log wage units: see equation (1).

Table 4: Decadal changes in employment shares and returns to education

	1990s	2000s	2010s
Share N.			
Non-graduates	0.638	0.584	0.588
Non-STEM graduates	0.297	0.329	0.321
STEM graduates	0.065	0.087	0.091
Return to education			
Non-STEM graduates v non-graduates			
Log salary	0.198	0.235	0.280
Worker efect	0.155	0.193	0.219
Firm effect	0.023	0.048	0.061
STEM graduates v non-graduates			
Log salary	0.286	0.561	0.740
Worker efect	0.148	0.377	0.505
Firm effect	0.095	0.180	0.222
Additional statistics			
Var. firm effect	0.029	0.037	0.035

Notes: This table reports key outcomes of interest, separately for the 1990s, 2000s and 2010s. The first rows show aggregate employment shares of each education group. We then report mean wage differentials between STEM/non-STEM graduates and non-graduates, which we disaggregate into AKM worker and firm effects, in line with equation (19). For this exercise, we estimate a separate AKM model for each decade. Finally, we report the variance of AKM firm effects by decade, corrected for measurement error using a split-sample procedure (see text).

Table 5: Regional effects on firm pay dispersion and sorting

	Mean: Firm AKM		Var: Firm AKM		Corr: Worker, Firm	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Base specifications						
Graduate share	0.321	0.378	0.055	0.062	0.693	0.617
	(0.054)	(0.071)	(0.006)	(0.018)	(0.133)	(0.287)
Panel B. Employment coefficients						
Log employment	0.011	0.009	0.002	0.000	0.028	0.035
	(0.003)	(0.015)	(0.000)	(0.003)	(0.007)	(0.049)
Panel C. Controlling for both						
Graduate share	0.321	0.382	0.058	0.062	0.577	0.629
	(0.054)	(0.063)	(0.009)	(0.018)	(0.184)	(0.269)
Log employment	0.000	0.011	0.000	0.000	0.001	0.04
	(0.003)	(0.009)	(0.000)	(0.002)	(0.010)	(0.047)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	No	Yes	No	Yes	No	Yes
Sample	98	98	98	98	98	98

Notes: Panel A shows the relationship between regional graduate share and local labor market outcomes. Panel B reproduces these estimates, but with log regional employment instead of graduate share. Panel C controls for both variables simultaneously. Odd-numbered columns exploit cross-sectional variation across regions, using equation (21). Even-numbered columns control for region fixed effects, as in equation (22), and so rely on within-region changes for identification. The dependent variables are the mean firm AKM premia (columns 1-2), the variance of firm AKM premia (columns 3-4), and the correlation between firm and worker AKM premia (columns 5-6). Sample consists of 49 regions observed in both 1995 and 2008 census years, for a total of 98 region-year observations. Observations are weighted by regional employment shares. Standard errors, clustered by region, in parentheses.

Table 6: Regional effects on education wage differentials

	Log Wage		Worker AKM		Firm AKM	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Non-STEM graduates v non-graduates						
Graduate share	0.823	0.326	0.583	0.069	0.188	0.191
	(0.112)	(0.240)	(0.090)	(0.202)	(0.027)	(0.062)
Panel B: STEM graduates v non-graduates						
Graduate share	1.784	1.138	1.062	0.48	0.576	0.480
	(0.246)	(0.572)	(0.209)	(0.439)	(0.051)	(0.139)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	No	Yes	No	Yes	No	Yes
Sample	98	98	98	98	98	98

Notes: Table shows the relationship between regional graduate share and local education wage differentials. Odd-numbered columns exploit cross-sectional variation across regions, using equation (21). Even-numbered columns control for region fixed effects, as in equation (22), and so rely on within-region changes for identification. Panel A explores wage differentials between non-STEM graduates and non-graduates, and Panel B between STEM graduates and non-graduates. The dependent variables are the mean log wage differential (columns 1-2), the mean differential in AKM worker effects (columns 3-4), and the mean differential in AKM firm effects (columns 5-6). Sample consists of 49 regions observed in both 1995 and 2008 census years, for a total of 98 region-year observations. Observations are weighted by regional employment shares. Standard errors, clustered by region, in parentheses.

Table 7: Spatial variation within firms

	Division effect			Degree share			Mean division effect
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean division effect	0.178	0.153		0.292	0.406		
	(0.035)	(0.041)		(0.066)	(0.062)		
Regional degree share			0.084			0.222	0.548
			(0.024)			(0.029)	(0.041)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Specification	OLS	IV	RF	OLS	IV	RF	FS
N	54,528	54,528	54,528	54,528	54,528	54,528	54,528

Notes: Columns 1-3 report estimates of β in equation (24), across firm $\times$ region (i.e. "divisional") observations, controlling for firm fixed effects. The dependent variable is the divisional wage premium, and the independent variable is the regional mean premium. IV specification (column 2) uses regional graduate share as an instrument, to correct for possible measurement error. Column 3 reports the reduced-form estimate. In columns 4-6, we redo the same specifications, but with divisional graduate share as the dependent variable. Column 7 shows the first-stage estimate. The sample consists of private-sector firms in 2012-2019. Observations are weighted by divisional employment. Standard errors, clustered by region, in parentheses.

Table 8: Differences between sectors in return to education

	Private sector	Public sector
Return to education		
Non-STEM graduates v non-graduates		
Log salary	0.280	0.309
Worker effect	0.218	0.300
Firm effect	0.063	-0.006
STEM graduates v non-graduates		
Log salary	0.739	0.504
Worker effect	0.490	0.470
Firm effect	0.234	0.021
Additional statistics		
Av. log wage	9.123	9.284
Av. worker effect	-0.047	0.152
Av. firm effect	0.014	-0.023
Var. firm effect	0.037	0.018

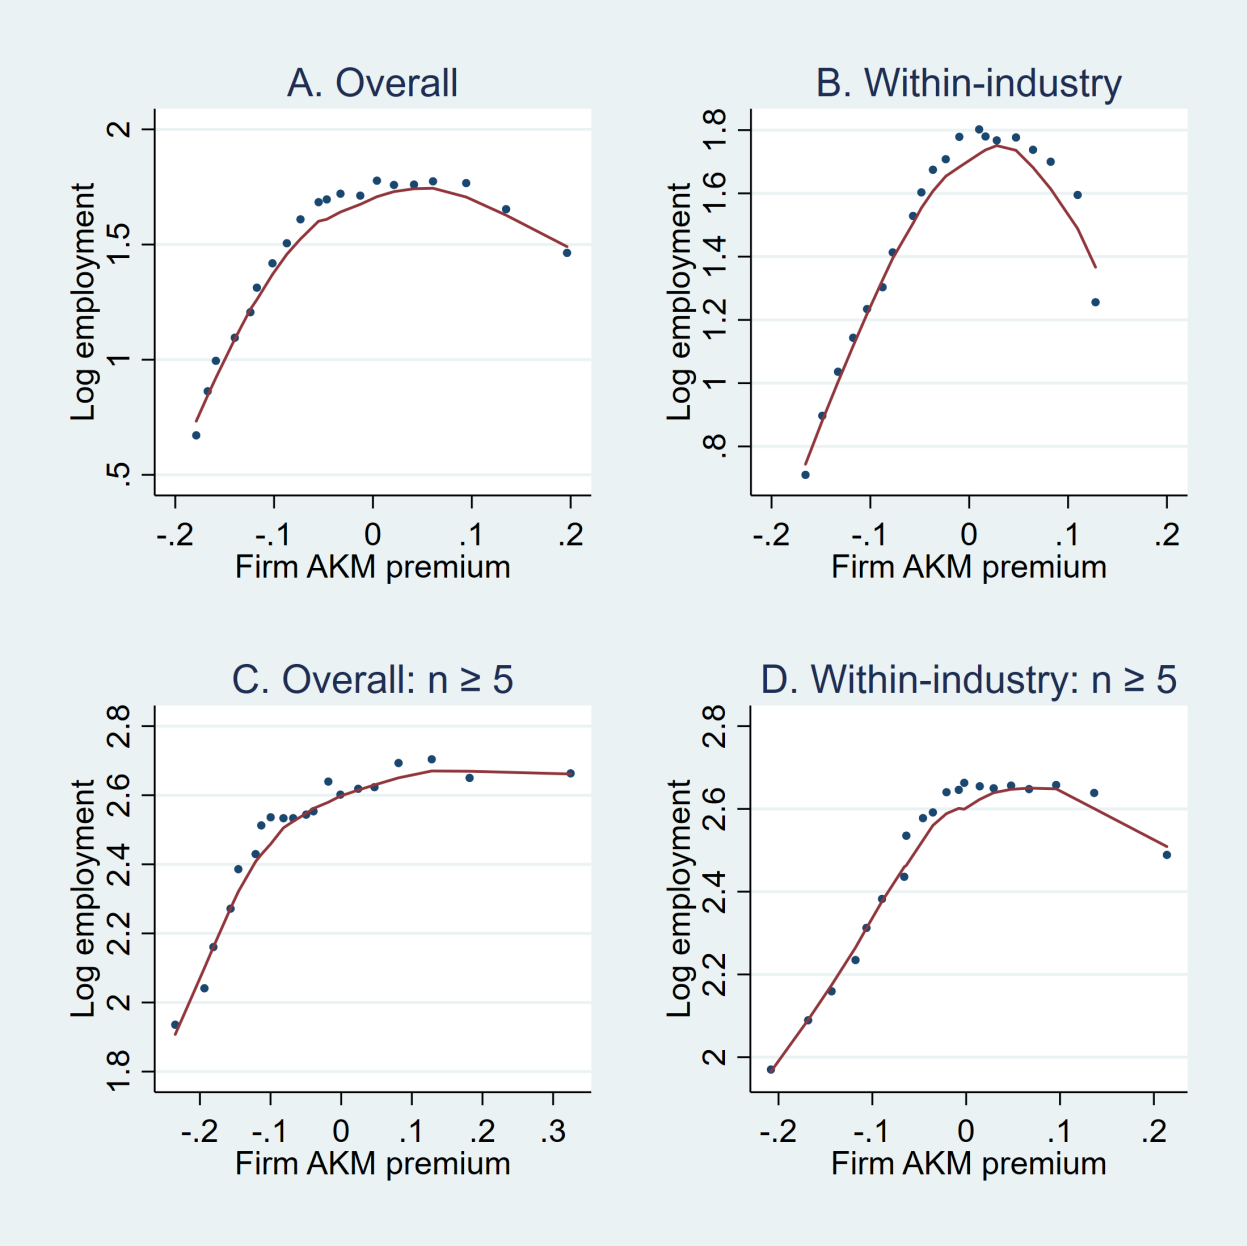
Notes: This table reports key outcomes of interest, separately for the private and public sectors for the period 2010-2019. The first rows show aggregate employment shares of each education group, by sector. We then report mean wage differentials between STEM/non-STEM graduates and non-graduates, which we disaggregate into AKM worker and firm effects, in line with equation (19). For this exercise, we estimate a common AKM model for both sectors. Finally, we report the variance of AKM firm effects by sector, corrected for measurement error using a split-sample procedure (see text).

Table 9: Regional effects on private and public sector outcomes

	Mean: Firm AKM		Var: Firm AKM		Corr: Worker, Firm	
	(1)	(2)	(3)	(4)	(5)	(6)
Graduate share	0.333	0.370	0.081	0.074	0.689	0.622
	(0.101)	(0.078)	(0.011)	(0.019)	(0.229)	(0.217)
Graduate share × Public	-0.334	-0.447	-0.067	-0.144	-1.255	-0.295
	(0.119)	(0.153)	(0.031)	(0.049)	(0.320)	(0.491)
	Mean: Worker AKM		Non-STEM grads		STEM grads	
	(7)	(8)	(9)	(10)	(11)	(12)
Graduate share	1.548	0.588	0.849	0.620	0.425	0.451
	(0.365)	(0.129)	(0.132)	(0.070)	(0.084)	(0.060)
Graduate share × Public	-0.569	-0.246	-0.729	-0.510	-0.265	-0.365
	(0.410)	(0.292)	(0.250)	(0.223)	(0.111)	(0.083)
Sector-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector-Region FE	No	Yes	No	Yes	No	Yes
Sample	196	196	196	196	196	196

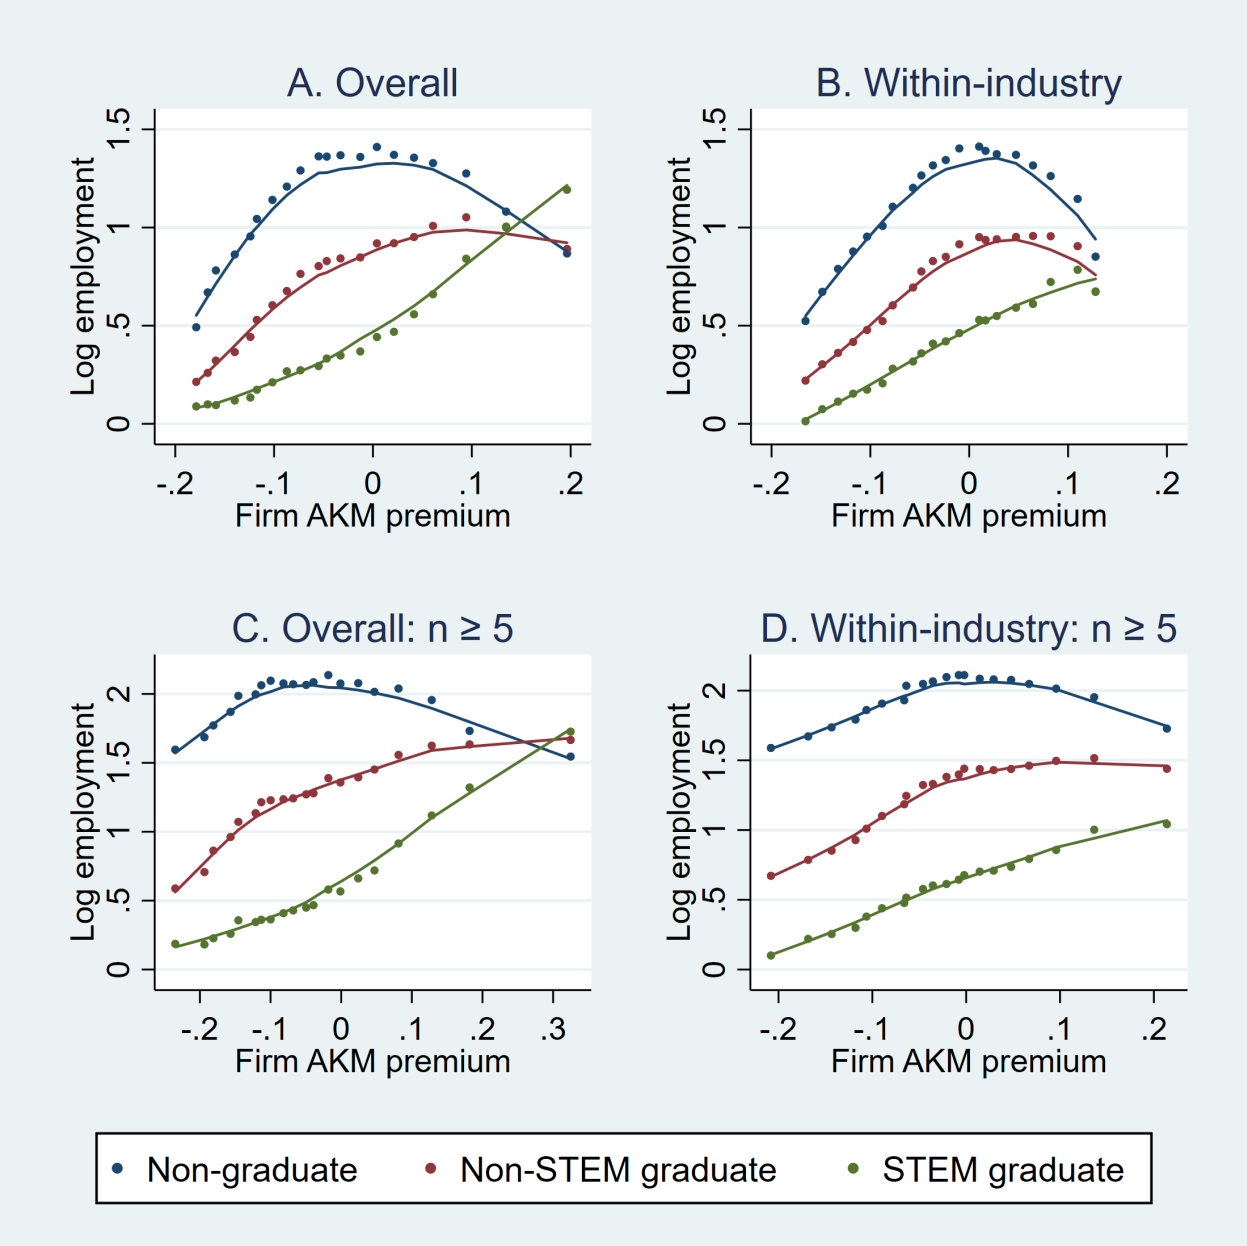
Notes: This table estimates the differential effects of regional graduate share (in the private sector) on private and public sector outcomes. Odd-numbered columns control for sector-year interacted fixed effects, in line with equation (26). Even-numbered columns additionally control for sector-region fixed effects, as in equation (27), and so rely on changes within sector-region cells for identification. The dependent variables are the mean firm AKM premia (columns 1-2), the variance of firm AKM premia (columns 3-4), and the correlation between firm and worker AKM premia (columns 5-6). The remaining columns focus on measures of worker quality: mean worker AKM, non-STEM graduate share, and STEM share. The outcomes in columns 3-6 are corrected for measurement error using a split-sample procedure (see text). Sample consists of two sectors (private and public) in 49 regions, observed in both 1995 and 2008 census years, for a total of 196 sector-region-year observations. Observations are weighted by employment shares of sector-region cells. Standard errors, clustered by region, in parentheses.

Figure 1: Employment by firm pay premium



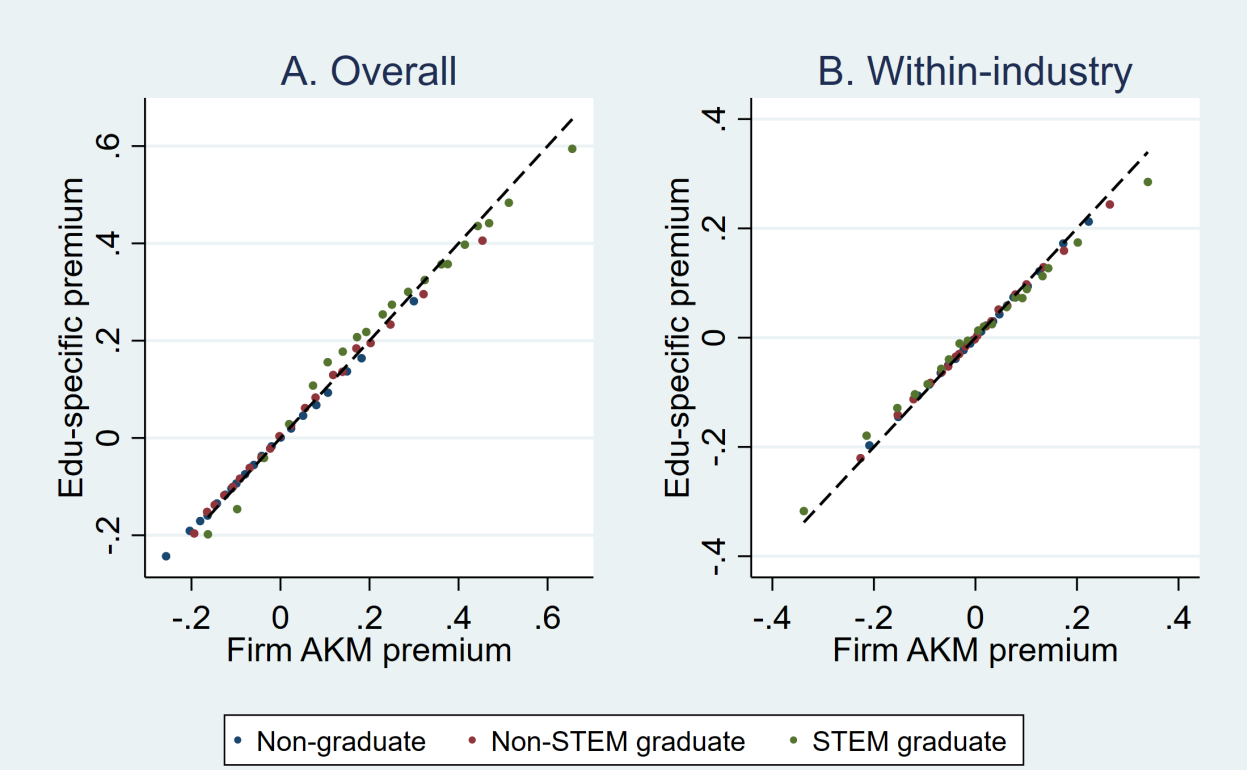
Notes: Panel A shows mean log firm employment across 20 bins (with equal numbers of firms), arranged by AKM firm premia. Firm premia are normalized to the worker-weighted mean. We implement a split-sample procedure to correct for measurement error in the firm premia, as described in Section 4.1. In Panel B, we remove industry fixed effects from both the y-variable (log firm employment) and the x-variable (firm premia). Panels C and D repeat this exercise after excluding firms with fewer than 5 employees. Sample consists of private-sector firms in 2010-2019.

Figure 2: Employment by education and firm pay premium



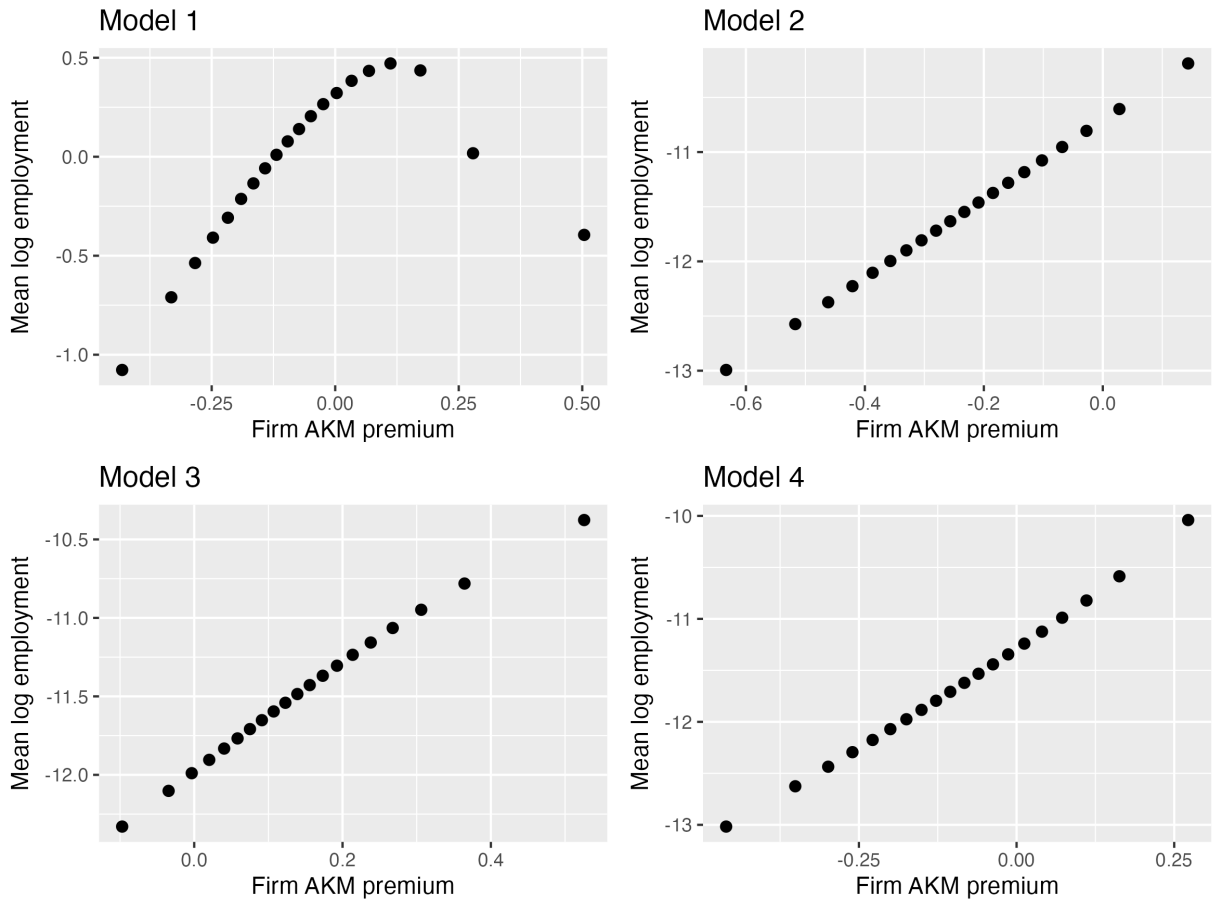
Notes: These plots repeat the exercise of Figure 1, but now showing mean log firm employment separately for three education groups: non-graduates, non-STEM graduates and STEM graduates. Sample consists of private-sector firms in 2010-2019.

Figure 3: Education-specific AKM firm premia



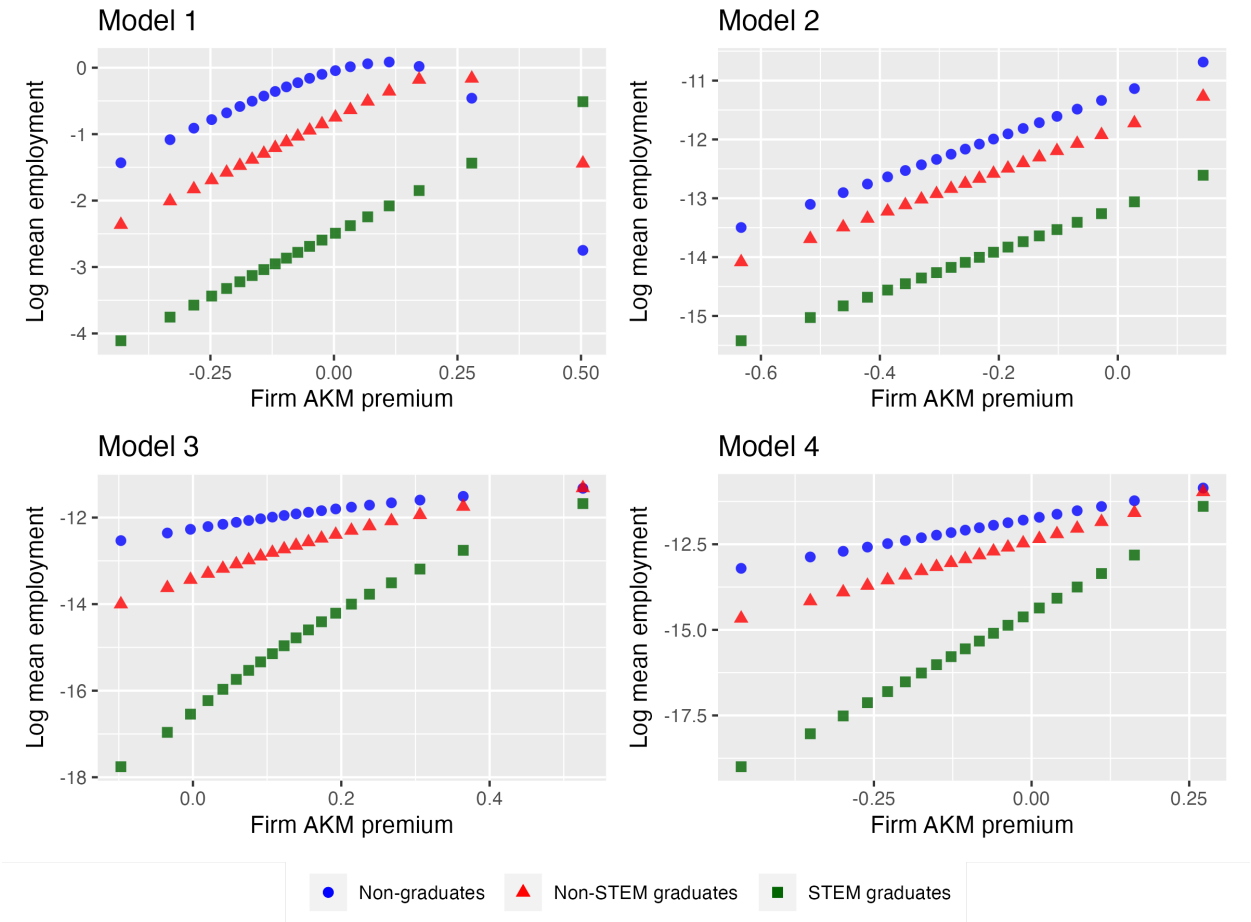
Notes: For this figure, we estimate AKM firm premia separately by education group, and plot these group-specific premia against the aggregate (i.e., full sample) firm premia, across 20 bins (ordered by the aggregate premia). The bins are defined separately by group, and contain equal numbers of group-specific workers. Group-specific and aggregate premia are normalized to zero for firms with mean (employment-weighted) aggregate premia. If wages are log-additive, the group-specific premia should line up perfectly on the 45 degree (dashed) line. Panel B repeats this exercise, after removing industry effects from the group-specific and aggregate premia. We implement a split-sample procedure to correct for measurement error in the firm premia, as described in Section 4.2. Sample includes private-sector firms in 2010-2019.

Figure 4: Employment by firm pay premium: Models



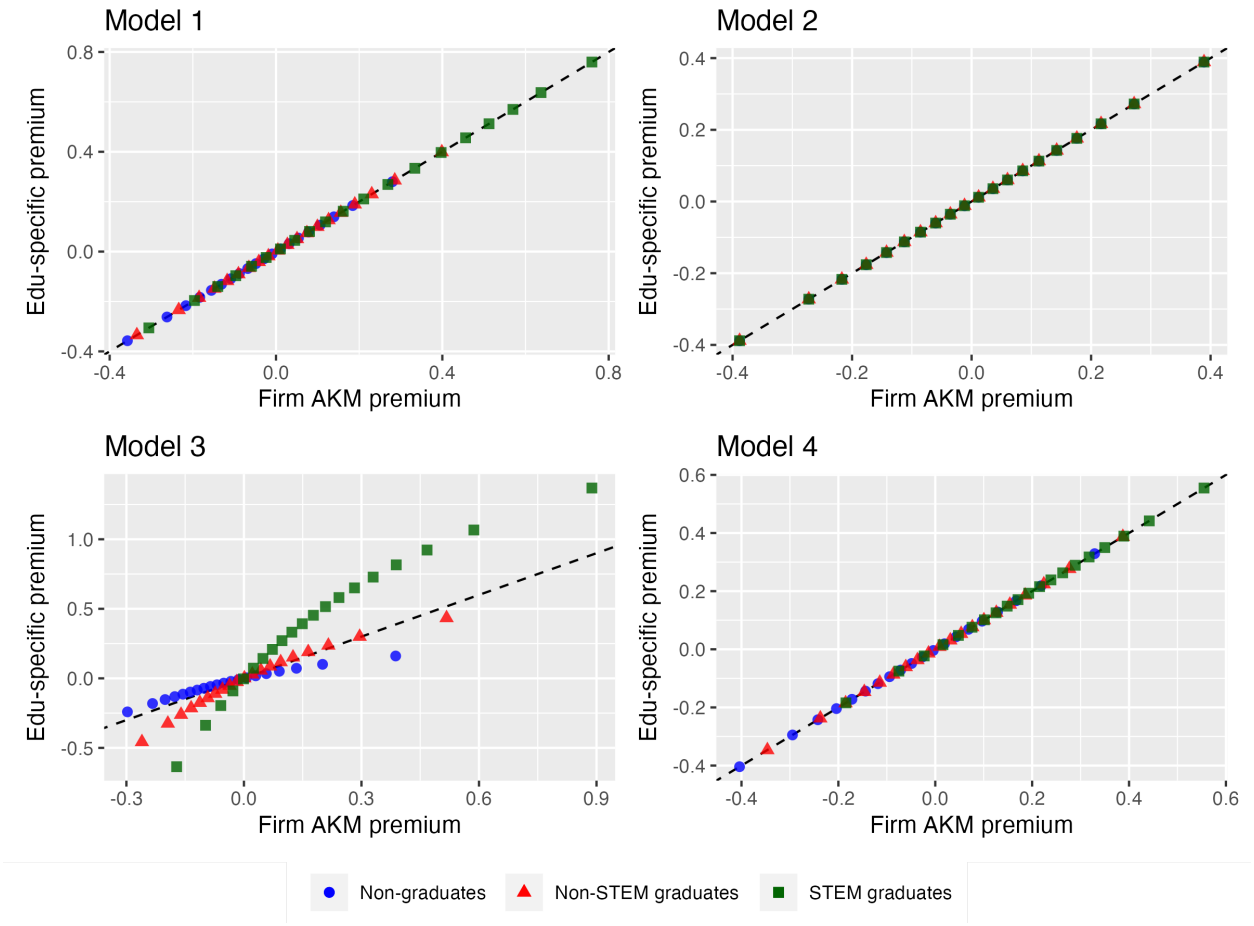
Notes: Figure shows the relationship between log firm employment and firm AKM wage premium for the four models described in the text. Each point represents one of 20 equally sized bins of firm wage premium. Employment is measured as mean log employment across firms within each bin.

Figure 5: Employment by education and firm pay premium: Models



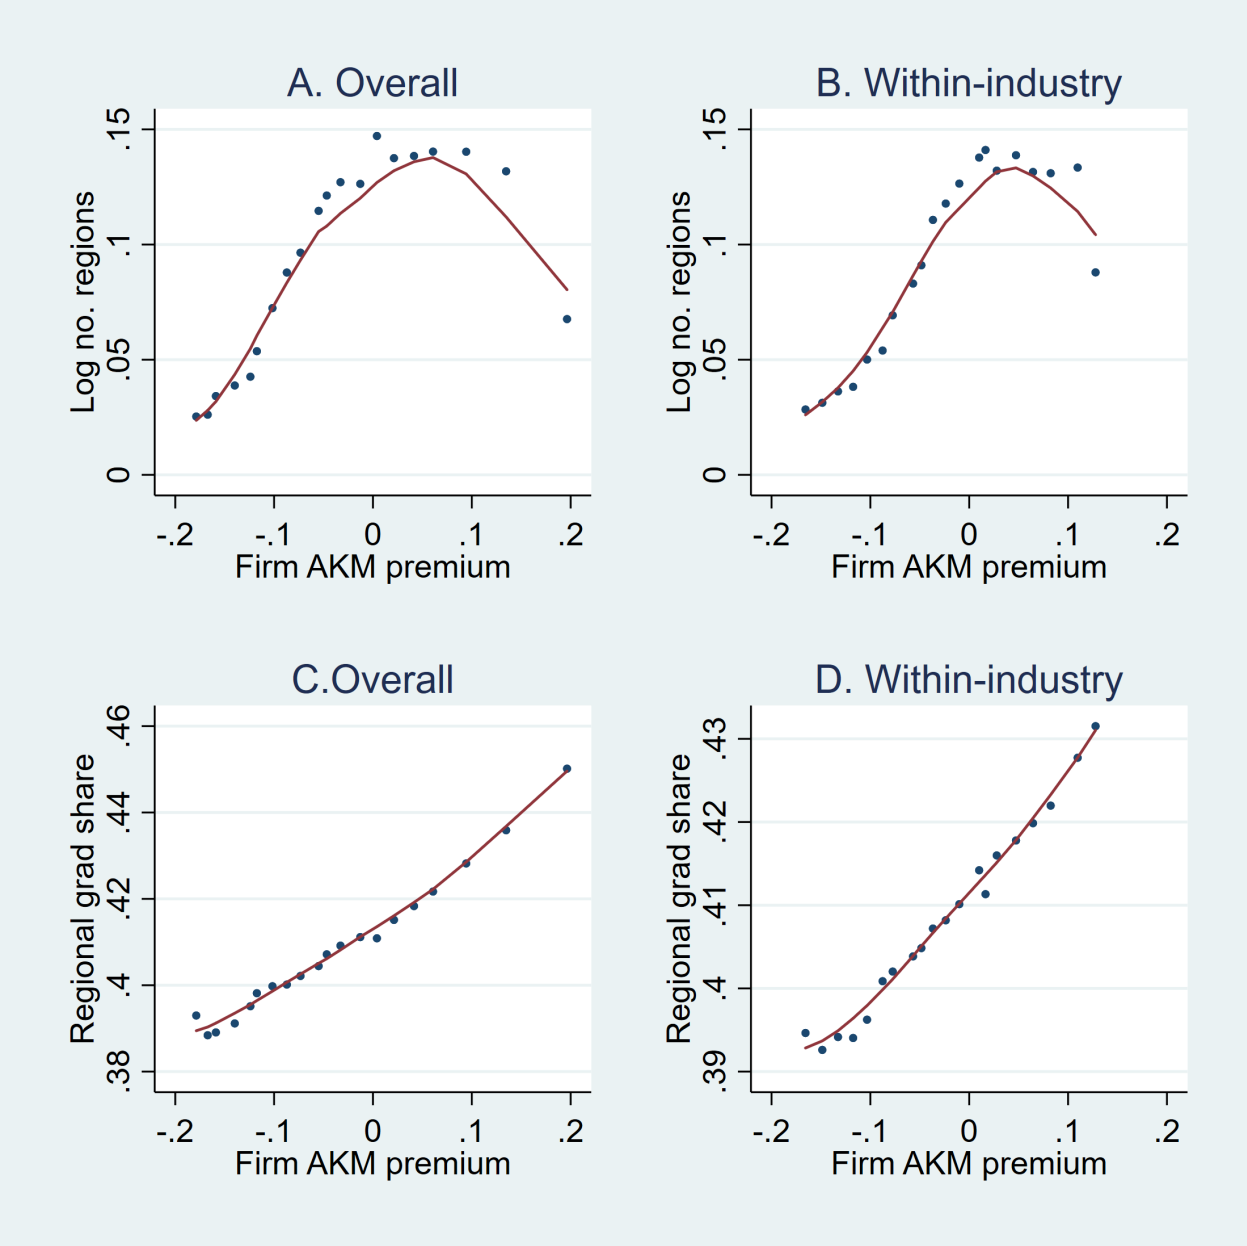
Notes: Figure shows the relationship between log employment by education group and firm AKM wage premium for the four models described in the text. Each point represents one of 20 equally sized bins of firm wage premium. Employment is measured as mean log employment within education group across firms in each bin.

Figure 6: Education-specific AKM firm premia: Models



Notes: The figure shows education-specific firm wage premia against overall firm AKM premium for the four models described in the text. Each point represents one of 20 bins, defined separately by group, and contain equal numbers of group-specific workers.

Figure 7: Firm location choices



Notes: Panel A shows the mean log number of regions in which firms operate across 20 bins (with equal numbers of firms), arranged by AKM firm premia. Firm premia are normalized to the worker-weighted mean. We implement a split-sample procedure to correct for measurement error in the firm premia, as described in Section 4.1. In Panel B, we remove industry fixed effects from both the y-variable (log number of regions) and the x-variable (firm premia). Panels C and D repeat this exercise, but the outcome is now the mean local graduate share in firms' regions of operation. Sample includes private sector firms in 2010-2019.

Appendices: For online publication

A Appendix tables and figures

Table A1: Calibrated Parameters Across Models

Parameter	Baseline	Skill-Neutral	Skill-Biased	Varying Elasticities
Labor supply parameters				
Labor supply elasticity (ε)	3.78	3.61	3.01	–
STEM-degree (ε_h)	–	–	–	10.16
Non-STEM (ε_m)	–	–	–	5.02
No degree (ε_l)	–	–	–	3.20
Productivity parameters				
Firm productivity variance (ν)	0.024	0.035	0.234	0.032
Non-STEM ratio ($\frac{p_m}{p_h}$)	0.50	-0.40	-0.16	-0.15
No-degree ratio ($\frac{p_l}{p_h}$)	0.31	-0.74	-0.07	-0.34
Non-STEM sensitivity (θ_m)	–	–	0.45	–
Non-degree sensitivity (θ_l)	–	–	0.20	–
Equity constraint parameters				
Non-STEM pay ratio (ϕ_m)	0.79	–	–	–
Non-degree pay ratio (ϕ_l)	0.59	–	–	–

Notes: This table presents the calibrated parameter values for each model variant. The baseline model features equity constraints and productivity differentials. The skill-neutral model allows only for firm productivity differences. The skill-biased model allows productivity differences to vary by education group. The varying elasticities model allows labor supply elasticities to differ across education groups.

Table A2: Regional distribution of skill shares and wages

	1995		2008	
	Mean	SD	Mean	SD
Graduate share	0.389	0.045	0.494	0.066
Non-STEM graduate share	0.327	0.039	0.399	0.048
STEM graduate share	0.063	0.016	0.096	0.035
Log wage	9.004	0.123	9.112	0.151
Firm AKM	0.007	0.020	0.011	0.035
Worker AKM	0.054	0.113	0.052	0.123

Notes: This table presents regional means and standard deviations of key variables in the 1995 and 2008 census years. Graduate share is the local fraction of workers with college degrees, which we disaggregate into non-STEM and STEM. AKM effects are estimated using employment records for the corresponding periods: 1993-1997 and 2006-2010. The sample includes 49 regions in each census year.

Table A3: Regional effects of STEM and non-STEM employment shares

	Mean: Firm AKM		Var: Firm AKM		Corr: Worker, Firm	
	(1)	(2)	(3)	(4)	(5)	(6)
Non-STEM grad share	0.168	0.126	0.150	-0.035	0.239	-0.014
	(0.106)	(0.084)	(0.026)	(0.072)	(0.120)	(0.331)
STEM grad share	0.646	0.665	0.316	0.475	1.059	1.334
	(0.112)	(0.106)	(0.037)	(0.060)	(0.225)	(0.332)
Region FE	No	Yes	No	Yes	No	Yes
Sample	98	98	98	98	98	98

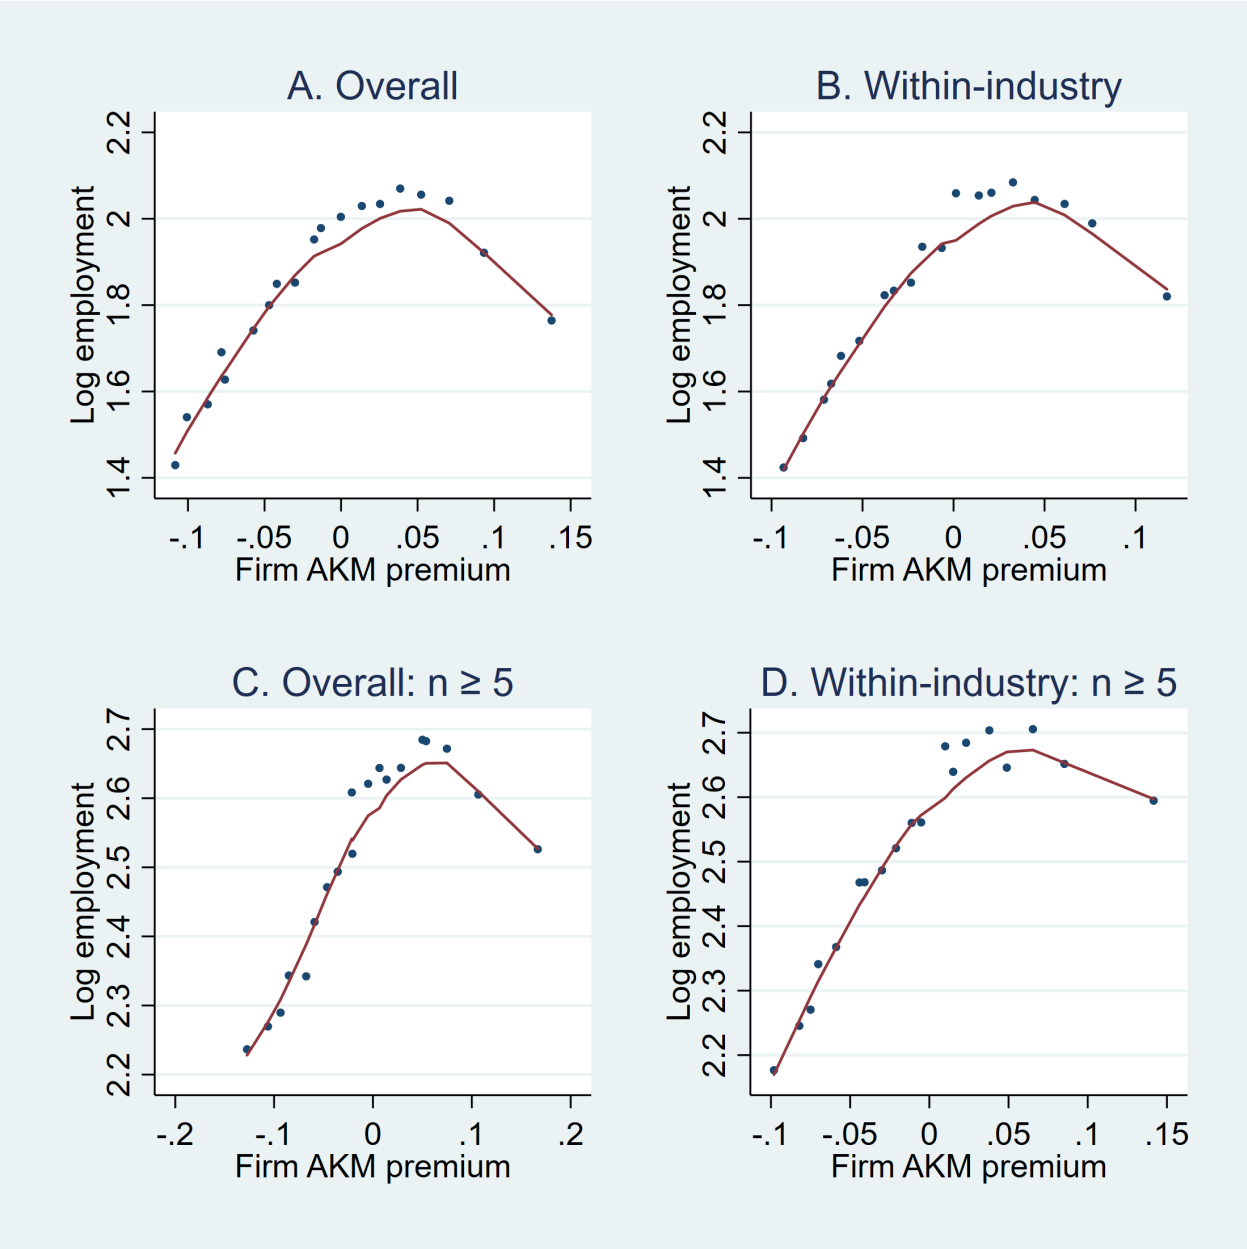
Notes: The table replicates Panel A of Table 5, but replacing the regional graduate share on the right-hand side with distinct STEM and non-STEM shares. Odd-numbered columns exploit cross-sectional variation across regions, using equation (21). Even-numbered columns control for region fixed effects, as in equation (22), and so rely on within-region changes for identification. The dependent variables are the mean firm AKM premia (columns 1-2), the variance of firm AKM premia (columns 3-4), and the correlation between firm and worker AKM premia (columns 5-6). Sample consists of 49 regions observed in both 1995 and 2008 census years, for a total of 98 region-year observations. Observations are weighted by regional employment shares. Standard errors, clustered by region, in parentheses.

Table A4: Regional variation in education wage differentials

	Log Wage		Worker AKM		Firm AKM	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Non-STEM graduates v non-graduates						
Non-STEM grad share	0.889	-0.370	0.675	-0.399	0.154	-0.016
	(0.180)	(0.215)	(0.130)	(0.234)	(0.042)	(0.072)
STEM grad share	0.680	1.116	0.386	0.600	0.259	0.425
	(0.296)	(0.302)	(0.234)	(0.272)	(0.071)	(0.061)
Panel B: STEM graduates v non-graduates						
Non-STEM grad share	1.329	-0.678	0.704	-0.751	0.485	-0.089
	(0.355)	(0.583)	(0.271)	(0.483)	(0.079)	(0.146)
STEM grad share	2.755	3.196	1.825	1.875	0.700	1.125
	(0.368)	(0.557)	(0.287)	(0.466)	(0.100)	(0.150)
Region FE	No	Yes	No	Yes	No	Yes
Sample	98	98	98	98	98	98

Notes: Table shows the relationship between regional graduate share and local education wage differentials. Odd-numbered columns exploit cross-sectional variation across regions, using equation (21). Even-numbered columns control for region fixed effects, as in equation (22), and so rely on within-region changes for identification. Panel A explores wage differentials between non-STEM graduates and non-graduates, and Panel B between STEM graduates and non-graduates. The dependent variables are the mean log wage differential (columns 1-2),the mean differential in AKM worker effects (columns 3-4), and the mean differential in AKM firm effects (columns 5-6). Sample consists of 49 regions observed in both 1995 and 2008 census years, for a total of 98 region-year observations. Observations are weighted by regional employment shares. Standard errors, clustered by region, in parentheses.

Figure A1: Employment by firm pay premium in Veneto



Notes: Panel A shows mean log firm employment across 20 bins (with equal numbers of firms) in the the Veneto Worker History (VWH) database, arranged by AKM firm premia. Firm premia are normalized to the worker-weighted mean. We implement a split-sample procedure to correct for measurement error in the firm premia, as described in Section 4.1. In Panel B, we remove industry fixed effects from both the y-variable (log firm employment) and the x-variable (firm premia). Panels C and D repeat this exercise after excluding firms with fewer than 5 employees. Sample consists of private sector firms in 1992-2001.

B Theoretical proofs for baseline model

B.1 Derivation of optimal inclusive wages (14) and (15)

Suppose the equity constraint binds, i.e. $\phi > \frac{p_l}{p_h}$. For inclusive firms, the l-type wage w_l will then equal ϕw_h; and the labor supply constraints will bind: i.e. $l_s = l_s(w_s)$ for $s = \{h, l\}$. We can then re-write the firm's problem in (4) as:

$$\max_{w_h} \pi(w_h) = (p_h - w_h) l_h(w_h) + (p_l - \phi w_h) l_l(\phi w_h) \tag{B1}$$

The first order condition for the h-type wage w_h is then:

$$(p_h - w_h) l_h'(w_h) + \phi(p_l - \phi w_h) l_l'(\phi w_h) = l_h(w_h) + \phi l_l(\phi w_h) \tag{B2}$$

After replacing $l_s(w_s)$ with (2), and using $w_s^* = \frac{\varepsilon}{1+\varepsilon} p_h$ from (10), and $\beta = \phi \frac{p_h}{p_l}$ from (12), we have:

$$w_h^I = \frac{1 + \frac{1}{\beta} \cdot \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}}{1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}} \cdot w_h^* \tag{B3}$$

which delivers (14); and (15) then follows from the binding pay constraint $w_l = \phi w_h$.

B.2 Derivation of optimal selective wage

For selective firms, only the h-type labor supply constraint binds, i.e. $l_h = l_h(w_h)$. We can then re-write the firm's problem in (4) as:

$$\max_{w_h, l_l} \pi(w_h) = (p_h - w_h) l_h(w_h) + (p_l - \phi w_h) l_l \tag{B4}$$

where l-type employment l_l is rationed, and must be strictly below the labor supply curve:

$$l_l < l_l(\phi w_h) \tag{B5}$$

Since marginal products are fixed by assumption, firms will only ration l_l if the l-type wage w_l (which is fixed by ϕw_h) exceeds their productivity p_l. But if this is indeed the case, firms will optimally reject all l-type workers: i.e., $l_l = 0$. More generally though, if the marginal product p_l is decreasing in l_l (e.g., if skill types are imperfect substitutes, or if there are diminishing returns to labor), the optimal l_l may lie between 0 and the labor supply curve.

Imposing $l_l = 0$, the first order condition for the h-type wage w_h is then:

$$(p_h - w_h)\, l_h'(w_h) = l_h(w_h) \tag{B6}$$

After replacing $l_h(w_h)$ with (2), we have:

$$w_h^S = \frac{\varepsilon}{1+\varepsilon} p_h = w_h^* \tag{B7}$$

where w_h^* is the optimal unconstrained wage.

B.3 Derivation of equilibrium equations (16) and (18)

Expressions for labor supply intercepts Ω_s

To solve for equilibrium, we first require expressions for the labor supply intercepts Ω_s, for $s = \{h, l\}$. Using equation (3), the intercept for h-type workers can be written as:

$$\Omega_h = \frac{n_h}{k}\left[\sigma\left(w_h^S\right)^{\varepsilon} + (1-\sigma)\left(w_h^I\right)^{\varepsilon}\right]^{-1} \tag{B8}$$

where n_h is the measure of h-type workers, and k is the measure of firms. The square brackets contain an average of the wages (with an ε exponent) of selective firms (weighted by the selective firm share σ) and inclusive firms (weighted $1-\sigma$). This weighted average represents the outside option of h-type workers.

Similarly, the intercept for l-type workers can be written as:

$$\Omega_l = \frac{n_l}{k}\left[(1-\sigma)\left(\phi w_h^I\right)^{\varepsilon}\right]^{-1} \tag{B9}$$

where n_l is the measure of l-type worker. Since l-type workers cannot access selective firms, the outside option in (B9) only accounts for inclusive firms.

Using the definitions of β and α in equations (12) and (17), the ratio of the two intercepts can be written as:

$$\frac{\Omega_l}{\Omega_h} = \frac{1-\alpha}{\alpha} \cdot \frac{\beta}{\phi^{1+\varepsilon}}\left[1 + \frac{\sigma}{1-\sigma}\left(\frac{w_h^S}{w_h^I}\right)^{\varepsilon}\right] \tag{B10}$$

Finally, replacing w_h^I and w_h^S with (B3) and (B7), we have:

$$\frac{\Omega_l}{\Omega_h} = \frac{1-\alpha}{\alpha} \cdot \frac{\beta}{\phi^{1+\varepsilon}}\left[1 + \frac{\sigma}{1-\sigma}\left(\frac{1+\phi^{1+\varepsilon}\frac{\Omega_l}{\Omega_h}}{1+\frac{\phi^{1+\varepsilon}}{\beta}\cdot\frac{\Omega_l}{\Omega_h}}\right)^{\varepsilon}\right] \tag{B11}$$

which is an equilibrium relationship between the intercept ratio $\frac{\Omega_l}{\Omega_h}$ and selective share σ.

To fix the equilibrium values of each, we need to assess the profits from the selective and inclusive strategies.

Expressions for inclusive and selective firm profits

Inserting the optimal inclusive wage (B3) into equation (4), and replacing $l_s(w_s)$ with (2), the inclusive profit can be written as:

$$\pi^I = \frac{\varepsilon^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \cdot \frac{\left(1 + \frac{\phi^{1+\varepsilon}}{\beta} \cdot \frac{\Omega_l}{\Omega_h}\right)^{1+\varepsilon}}{\left(1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}\right)^{\varepsilon}} \cdot \Omega_h p_h^{1+\varepsilon} \tag{B12}$$

Similarly, inserting the optimal selective wage (B7) into equation (B4), and replacing $l_h(w_h)$ with (2), the selective profit can be written as:

$$\pi^S = \frac{\varepsilon^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \Omega_h p_h^{1+\varepsilon} \tag{B13}$$

Equilibrium with zero workplace segregation: $\sigma = 0$

For an equilibrium with zero workplace segregation ($\sigma = 0$), firms must strictly prefer the inclusive strategy: i.e. $\pi^I > \pi^S$. Using (B12) and (B13), this implies:

$$\left(1 + \frac{\phi^{1+\varepsilon}}{\beta} \cdot \frac{\Omega_l}{\Omega_h}\right)^{1+\varepsilon} > \left(1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}\right)^{\varepsilon} \tag{B14}$$

But imposing $\sigma = 0$ on (B11), we have:

$$\frac{\Omega_l}{\Omega_h} = \frac{1-\alpha}{\alpha} \cdot \frac{\beta}{\phi^{1+\varepsilon}} \tag{B15}$$

And after inserting this equation (B14) and rearranging:

$$\beta < \frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}} - \alpha}{1-\alpha} \tag{B16}$$

which is the threshold condition for a $\sigma = 0$ equilibrium in equation (16).

Equilibrium with partial workplace segregation: $\sigma > 0$

For an equilibrium with partial workplace segregation ($\sigma > 0$), firms must be indifferent between the selective and inclusive strategies: i.e. $\pi^I = \pi^S$. Equating (B12) and (B13), this

implies:

$$\left(1+\frac{\phi^{1+\varepsilon}}{\beta}\cdot\frac{\Omega_l}{\Omega_h}\right)^{1+\varepsilon}=\left(1+\phi^{1+\varepsilon}\frac{\Omega_l}{\Omega_h}\right)^{\varepsilon} \tag{B17}$$

Imposing this on (B11) yields:

$$\frac{\Omega_l}{\Omega_h}=\frac{1-\alpha}{\alpha-\sigma}\cdot\frac{\beta}{\phi^{1+\varepsilon}} \tag{B18}$$

And replacing $\frac{\Omega_l}{\Omega_h}$ in equation (B17) with (B18):

$$\left(1+\frac{1-\alpha}{\alpha-\sigma}\right)^{1+\varepsilon}=\left(1+\beta\frac{1-\alpha}{\alpha-\sigma}\right)^{\varepsilon} \tag{B19}$$

This is an implicit equation which solves for $\tilde{\sigma}$ in equation (18), i.e. the value of σ in an equilibrium with partial workplace segregation.

B.4 Proof of Proposition 1c: Wage compression effects

Effect on expected log wages by skill type

In this appendix, we study the impact of the equity constraint on expected log wages by skill type and aggregate earnings. We begin with h-type wages. Let $E\left[\log w_h|\beta>1\right]$ denote the expected log wage of h-types in an economy with a binding equity constraint. This is a weighted average of log wages paid by selective and inclusive firms, with weights equal to their shares of h-type employment:

$$E\left[\log w_h|\beta>1\right]=\frac{(1-\sigma)\,l_h\left(w_h^I\right)\log w_h^I+\sigma l_h\left(w_h^S\right)\log w_h^S}{(1-\sigma)\,l_h\left(w_h^I\right)+\sigma l_h\left(w_h^S\right)} \tag{B20}$$

In a counterfactual unconstrained economy, all firms offer h-types the unconstrained optimum w_h^*, as defined by equation (10). As Appendix B.2 shows, the optimal selective wage w_h^S is equal to w_h^*. Using (B20), the impact of the equity constrained can then be written as:

$$E\left[\log w_h|\beta>1\right]-\log w_h^*=\frac{(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}}{(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}+\sigma}\log\frac{w_h^I}{w_h^*} \tag{B21}$$

Note the optimal inclusive wage w_h^I is less than the unconstrained optimum w_h^*: see equation (14). It follows that $E\left[\log w_h\right]-\log w_h^*$ is negative: i.e., the equity constraint reduces the expected log h-type wage.

We next turn to l-type wages. Let $E\left[\log w_l|\beta>1\right]$ denote the expected log wage of l-types in an economy with a binding equity constraint. Since l-types are denied employment

by selective firms, this is simply equal to the log inclusive wage. The impact of the pay constraint (relative to the unconstrained optimum) can then be written as:

$$E\left[\log w_l | \beta > 1\right] - \log w_l^* = \log \frac{w_l^I}{w_l^*} \tag{B22}$$

From equation (15), we know the inclusive wage w_l^I must exceed the unconstrained optimum w_l^*. So this expression must be positive: i.e., the equity constraint increases the expected log l-type wage.

Effect on aggregate earnings and profit

Finally, we show that aggregate earnings are unaffected by the equity constraint. Since output in this model is fixed by assumption (workers are equally productive at all firms), it is sufficient to show that profit is unaffected by the equity constraint.

We begin by solving for profit π^* in an unconstrained economy. Applying the optimal wage (10) to (3), the labor supply intercepts for skill type s will equal:

$$\Omega_s^* = \left(\frac{1+\varepsilon}{\varepsilon}\right)^{\varepsilon} \frac{n_s}{k} \cdot p_s^{-\varepsilon} \tag{B23}$$

Using this expression, and applying the binding labor supply curve (2) and optimal wage (10) to the profit function (4), we have:

$$\pi^* = \frac{1}{\alpha} \cdot \frac{n_h}{k} \cdot \frac{p_h}{1+\varepsilon} \tag{B24}$$

where the h-type output share α is defined by (17).

Next, we turn to profit under a binding equity constraint, which we denote $\pi|\beta > 1$. Since firms are identical (and earn equal profit), we can use the profit of inclusive firms from equation (B12):

$$\pi|\beta > 1 = \frac{\varepsilon^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}} \cdot \frac{\left(1 + \frac{\phi^{1+\varepsilon}}{\beta} \cdot \frac{\Omega_l}{\Omega_h}\right)^{1+\varepsilon}}{\left(1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}\right)^{\varepsilon}} \cdot \Omega_h p_h^{1+\varepsilon} \tag{B25}$$

Inserting expressions for the h-type labor supply intercept Ω_h from (B8), the equilibrium intercept ratio $\frac{\Omega_l}{\Omega_h}$ from (B18), the optimal unconstrained wage w_h^* from (10), and the optimal inclusive wage w_h^I from (14), this can be written as:

$$\pi|\beta > 1 = \frac{\left(1 + \frac{1-\alpha}{\alpha-\sigma}\right)^{1+\varepsilon}}{\sigma\left(1 + \beta\frac{1-\alpha}{\alpha-\sigma}\right)^{\varepsilon} + (1-\sigma)\left(1 + \frac{1-\alpha}{\alpha-\sigma}\right)^{\varepsilon}} \cdot \frac{n_h}{k} \cdot \frac{p_h}{1+\varepsilon} \tag{B26}$$

We now consider two cases. In a non-segregated equilibrium, the selective share σ will equal zero; and equation (B26) will collapse to the unconstrained profit π^* in (B24). In a partially segregated equilibrium, the equal profit condition in equation (B19) ensures that $\left(1+\beta\frac{1-\alpha}{\alpha-\sigma}\right)^{\varepsilon} = \left(1+\frac{1-\alpha}{\alpha-\sigma}\right)^{\varepsilon}$; so again, equation (B26) will collapse to the unconstrained profit π^*. Therefore, profit is unaffected by the equity constraint; so the same must be true of aggregate earnings.

B.5 Proof of Proposition 1d: Amenity and welfare effects

Let $\bar{u}_s$ denote the expected utility of skill type s workers, and let $\bar{a}_s$ denote their expected amenity match. Given the assumption that amenity effects are distributed type-1 extreme value, $\bar{u}_s$ will equal the log of the inclusive value:

$$\bar{u}_s = \log \int_f w_{sf}^{\varepsilon} df + \gamma \tag{B27}$$

where γ is Euler's constant. From equation (1), the expected amenity match $\bar{a}_s$ can then be imputed by subtracting ε times the expected log wage:

$$\bar{a}_s = \log \int_f w_{sf}^{\varepsilon} df - \varepsilon E\left[\log w_s\right] + \gamma \tag{B28}$$

Proposition 1d states that the equity constraint increases the expected match $\bar{a}_s$ for both skill types (relative to the unconstrained optimum), if the constraint has sufficient bite (such that the selective share σ exceeds zero). We prove this result for each skill type in turn.

Effect on expected amenity match for h-types

Let $\bar{a}_s^*$ denote the expected amenity match in an unconstrained economy. For h-types, wages are fixed at the unconstrained optimum w_h^* in all firms; so $\bar{a}_h^*$ is simply equal to Euler's constant γ.

Next, let $\bar{a}_h|\beta > 1$ denote the expected amenity match for h-types in an economy with a binding equity constraint. Using (B28) and (B21), this can be written as:

$$\bar{a}_h|\beta > 1 = \log\left[(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon} + \sigma\right] - \frac{(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}}{(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon} + \sigma} \log\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon} + \gamma \tag{B29}$$

The impact of the equity constraint, compared to an unconstrained counterfactual, is there-

fore:

$$(\bar{a}_h|\beta>1)-\bar{a}_h^*=\log\left[(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}+\sigma\right]-\frac{(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}}{(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}+\sigma}\log\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon} \tag{B30}$$

Since inclusive firms pay less than the unconstrained optimum w_h^*, the term $\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}$ must lie between 0 and 1. Notice that for $\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}=1$, $(\bar{a}_h|\beta>1)-\bar{a}_h^*$ is equal to zero. But after differentiating (B30) with respect to $\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}$, it can be shown that $(\bar{a}_h|\beta>1)-\bar{a}_h^*$ is strictly increasing in $\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}$ for $\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon}<1$, as long as the selective share σ exceeds zero. It follows that $(\bar{a}_h|\beta>1)-\bar{a}_h^*$ must be less than zero, for $\sigma>0$: i.e. an equity constraint with sufficient bite (such that $\sigma>0$) reduces the expected amenity match for h-types.

Since the equity constraint also reduces the expected log wage of h-types (Proposition 1c), we can conclude that their expected utility (which is the sum of the two) must decrease.

Effect on expected amenity match for l-types

Let $\bar{a}_l|\beta>1$ denote the expected amenity match for l-type in an economy with a binding equity constraint. Using (B28) and (B22), this can be written as:

$$\bar{a}_l|\beta>1=\log(1-\sigma)\,w_l^I-\log w_l^I+\gamma \tag{B31}$$

The impact of the equity constraint, compared to an unconstrained counterfactual, is therefore:

$$(\bar{a}_l|\beta>1)-\bar{a}_l^*=\log(1-\sigma) \tag{B32}$$

which is less than zero, if the selective share σ exceeds zero. That is, the equity constraint with sufficient bite (such that $\sigma>0$) reduces the expected amenity match for l-types.

This reduction in the expected amenity match offsets the increase in the expected log wage of l-types (from Proposition 1c). And it turns out the overall impact on their expected utility is ambiguous. To see this, consider two extreme cases. (i) Suppose the bite of the equity constraint delivers an equilibrium selective share σ of zero. In this case, there is no impact on the expected amenity match in (B32), but the expected log wage in (B32) does increase; so expected utility must increase also. (ii) Suppose the h-type output share is arbitrarily close to 1, so the same must be true of the equilibrium selective share σ. This ensures an arbitrarily large negative effect on the expected amenity match in (B32), which will dominate the impact on the expected log wage in (B32). In this case, the equity constraint must reduce expected utility.

B.6 Proof of Proposition 2a: Negative firm size premium

In the baseline model with productively identical firms, Proposition 3 states that high-pay (i.e., selective) firms will have lower employment overall. To prove this result, we derive firm size for the selective and inclusive strategies.

Selective firms only employ h-type workers, and they pay the unconstrained optimal wage: i.e., $w_h^S = w_h^*$. Therefore, using the labor supply function (2), their firm size is equal to:

$$l_h\left(w_h^S\right) = \Omega_h \left(w_h^*\right)^\varepsilon \tag{B33}$$

where Ω_h is the h-type labor supply intercept.

Inclusive firms employ both h- and l-type workers: they pay the former w_h^I and the latter $w_l^I = \phi w_h^I$. Therefore, using the labor supply functions in (2), their firm size is equal to:

$$\begin{aligned} l_h\left(w_h^I\right) + l_l\left(w_l^I\right) &= \Omega_h\left(w_h^I\right)^\varepsilon + \Omega_l\left(\phi w_h^I\right)^\varepsilon \\ &= \left(1 + \frac{\Omega_l}{\Omega_h}\phi^\varepsilon\right)\Omega_h\left(w_h^I\right)^\varepsilon \end{aligned} \tag{B34}$$

Replacing the optimal inclusive wage w_h^I with equation (14) gives:

$$l_h\left(w_h^I\right) + l_l\left(w_l^I\right) = \left(1 + \frac{\Omega_l}{\Omega_h}\phi^\varepsilon\right)\left(\frac{1 + \frac{1}{\beta}\cdot\phi^{1+\varepsilon}\frac{\Omega_l}{\Omega_h}}{1 + \phi^{1+\varepsilon}\frac{\Omega_l}{\Omega_h}}\right)^\varepsilon \Omega_h\left(w_h^*\right)^\varepsilon \tag{B35}$$

Since we are comparing selective and inclusive firms, we must be in an equilibrium with positive selective share ($\sigma > 0$). Therefore, the equilibrium ratio $\frac{\Omega_l}{\Omega_h}$ of the labor supply intercepts can be summarized by equation (B18). After replacing $\frac{\Omega_l}{\Omega_h}$ with (B18) and rearranging, we have:

$$l_h\left(w_h^I\right) + l_l\left(w_l^I\right) = \frac{1-\sigma}{\alpha-\sigma}\cdot\frac{\alpha-\sigma+(1-\alpha)\frac{\beta}{\phi}}{\alpha-\sigma+(1-\alpha)\beta}\cdot\Omega_h\left(w_h^*\right)^\varepsilon \tag{B36}$$

Since $\sigma < \alpha$ in equilibrium and $\phi \leq 1$, this expression must exceed the selective firm size (B33). This confirms that selective firms do indeed have lower employment overall.

B.7 Derivation of equation (19): Decomposition of skill differential

We begin by deriving a simple expression for the ratio of the inclusive wage to the optimal unconstrained wage, i.e. $\frac{w_h^I}{w_h^*}$. Replacing the intercept ratio $\frac{\Omega_l}{\Omega_h}$ with (B18) in equation (14), we have:

$$\frac{w_h^I}{w_h^*} = \frac{1 + \frac{1-\alpha}{\alpha-\sigma}}{1 + \beta\frac{1-\alpha}{\alpha-\sigma}}\cdot w_h^* \tag{B37}$$

Using (18), this can be re-written as:

$$\frac{w_h^I}{w_h^*} = \left(\frac{\alpha - \sigma}{1 - \sigma}\right)^{\frac{1}{\varepsilon}} \tag{B38}$$

Next, we turn to the skill differential. The expected log h-type wage is given by equation (B20). And since l-types are denied access to selective firms, they all receive the inclusive firm wage: i.e. $w_l^I = \phi w_h^I$. Subtracting one from the other, the expected skill differential is:

$$E\left[\log w_h\right] - E\left[\log w_l\right] = \frac{(1-\sigma)\, l_h\left(w_h^I\right) \log w_h^I + \sigma l_h\left(w_h^S\right) \log w_h^S}{(1-\sigma)\, l_h\left(w_h^I\right) + \sigma l_h\left(w_h^S\right)} - \log \phi w_h^I \tag{B39}$$

Applying the labor supply function (2), and given that the selective wage w_h^S is equal to the optimal unconstrained wage w_h^*, we have:

$$E\left[\log w_h\right] - E\left[\log w_l\right] = \frac{\sigma}{(1-\sigma)\left(\frac{w_h^I}{w_h^*}\right)^{\varepsilon} + \sigma} \log \frac{w_h^*}{w_h^I} - \log \phi \tag{B40}$$

And after replacing $\frac{w_h^I}{w_h^*}$ with (B38), and ϕ with (12), we reach equation (19) in the main text.

B.8 Proof of Proposition 3: Impact of h-type output share α

Suppose first that the h-type output share α is sufficiently large, such that $\frac{1-\alpha}{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}} - \alpha} > \beta$; so the selective share σ exceeds zero: see equation (16). For the purposes of this proof, it is useful to define the function $\Lambda(\beta, \varepsilon)$ as the solution of the implicit equation:

$$(1 + \Lambda)^{1+\varepsilon} = (1 + \beta\Lambda)^{\varepsilon} \tag{B41}$$

This is identical to the equilibrium equation (B19), except with $\frac{1-\alpha}{\alpha-\sigma}$ replaced by Λ, which exceeds zero if $\sigma > 0$. Using this definition, we can express equilibrium as:

$$\Lambda(\beta, \varepsilon) = \frac{1 - \alpha}{\alpha - \sigma} \tag{B42}$$

From equation (B42), since Λ is fixed by the exogenous parameters β and ε (and invariant to α), the selective share σ must then be increasing in α.

Next, consider the between-firm component in equation (19), which is equal to $\frac{\sigma}{\alpha} \log\left(\frac{1-\sigma}{\alpha-\sigma}\right)^{\frac{1}{\varepsilon}}$.

Using (B42), this can be re-written as:

$$\text{Between-firm} = \left[1 - \frac{1-\alpha}{\alpha} \cdot \frac{1}{\Lambda(\beta,\varepsilon)}\right] \log\left(1 + \Lambda(\beta,\varepsilon)\right)^{\frac{1}{\varepsilon}} \tag{B43}$$

Holding the exogenous parameters β and ε fixed, the between-firm component in (B43) must be increasing in α. This proves the first part of the proposition.

Now suppose that $\frac{1-\alpha}{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}}-\alpha} < \beta$. Equation (16) then shows that the selective share σ is fixed at zero and unaffected by α. We are therefore in a zero-segregation equilibrium, where all firms pursue the inclusive hiring strategy and offer identical wages: w_h^I for h-types, and w_l^I for l-types. Replacing the intercept ratio $\frac{\Omega_l}{\Omega_h}$ with (B15) in equation (14), the h-type inclusive wage in a zero-segregation equilibrium is equal to:

$$w_h^I = \frac{1}{\alpha + (1-\alpha)\beta} \cdot w_h^* \tag{B44}$$

which is increasing in α (as well as in h-type productivity p_h, via the optimal unconstrained wage term w_h^*). At the same time, the l-type wage is equal to ϕw_h^I (where ϕ is fixed by the equity constraint), so any productive benefits of larger α are shared equally with l-types.

B.9 Implications of incomplete pass-through to equity constraint ϕ

In Section 2.5, we study how changes in relative h-type productivity $\frac{p_h}{p_l}$ and relative labor supply $\frac{n_h}{n_l}$ affect skill wage differentials, holding the constraint bite β fixed. We show that both the within-firm and between-firm components of skill differentials (in equation (19)) are increasing in relative productivity; and the latter is also increasing in relative labor supply.

Given the definition of β in equation (12), a fixed β implies perfect pass-through of skill productivity differentials $\frac{p_h}{p_l}$ to within-firm pay differentials (and the equity constraint) ϕ. In this appendix, we consider the implications of incomplete pass-through. To be more precise, suppose the equity constraint ϕ takes the form: $\phi = \phi_0 \left(\frac{p_l}{p_h}\right)^{\phi_1}$. If the elasticity ϕ_1 is equal to 1, we have perfect pass-through to within-firm pay differentials; so the constraint bite $\beta \equiv \phi \frac{p_h}{p_l}$ is fixed at ϕ_0. But if $\phi_1 < 1$, we have imperfect pass-through; and the constraint bite β will be increasing in relative h-type productivity.

This will reduce the effect of relative productivity $\frac{p_h}{p_l}$ on *within-firm* skill differentials (in equation (19)), but increase its effect on the *between-firm* component. Regarding the latter, we know from Proposition 3 that the selective share σ is increasing in the h-type output share α (and hence their relative productivity), holding the constraint bite β fixed. But if pass-through is only partial, skill-biased productivity growth will increase the bite β; and for

given α, Proposition 1 shows this will increase the equilibrium selective share σ. It follows that imperfect pass-through will amplify the original (positive) effect of α on σ and hence on the between-firm component. Intuitively, to the extent that firms cannot differentiate pay within firms, the quantity-quality trade-off becomes more acute; and workplace segregation increases in its stead.

C Extension with heterogeneous firms

In this appendix, we extend our baseline model to account for skill-neutral heterogeneity in firm productivity. In a given firm f with firm-specific parameter x_f, suppose the h-type and l-type marginal products are equal to $p_{hf} = x_f p_h$ and $p_{lf} = x_f p_l$ respectively, where $\tilde{x}_f \equiv \log x_f$ has distribution F across firms, where F is normal with mean 0 and variance ν. For the purposes of this analysis, suppose the equity constraint binds, and β exceeds $\frac{\left(\frac{1}{\alpha}\right)^{\frac{1}{\varepsilon}}-\alpha}{1-\alpha}$; so the equilibrium selective share σ exceeds 0.

C.1 Characterization of equilibrium with heterogeneous firms

We begin by characterizing the equilibrium in this extended model. Building from equation (10), for a firm with productivity x, the unconstrained optimum wage for skill type $s = \{h, l\}$ can be written as:

$$w_s^*(x) = \frac{\varepsilon}{1+\varepsilon} p_s x \tag{C1}$$

Selective firms with productivity x pay the unconstrained optimum to h-types:

$$w_h^S(x) = w_h^*(x) \tag{C2}$$

Replacing w_h^* with $w_h^*(x)$ in equation (14), inclusive firms with productivity x offer a wage equal to:

$$w_h^I(x) = \frac{1 + \frac{1}{\beta} \cdot \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}}{1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}} w_h^*(x) \tag{C3}$$

Replacing p_h with $p_h x$ in equations (B12) and (B13), the profits associated with these strategies are:

$$\pi^S(x) = \frac{\varepsilon^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \Omega_h (p_h x)^{1+\varepsilon} \tag{C4}$$

and

$$\pi^I(x) = \frac{\varepsilon^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \cdot \frac{\left(1 + \frac{\phi^{1+\varepsilon}}{\beta} \cdot \frac{\Omega_l}{\Omega_h}\right)^{1+\varepsilon}}{\left(1 + \phi^{1+\varepsilon} \frac{\Omega_l}{\Omega_h}\right)^{\varepsilon}} \cdot \Omega_h (p_h x)^{1+\varepsilon} \tag{C5}$$

Comparing (B12) and (B13), it is clear that the productivity parameter x makes no difference to the relative profits of the two strategies; and hence, x does not affect the choice of strategy. It follows that selective and inclusive firms will be distributed identically in terms of x.

This allows us to characterize the pay distributions among selective and inclusive firms. Let F^s be the distribution of log h-type wages among selective firms, i.e. $\tilde{w}_h^S \sim F^S$, where the tilde indicates a log variable: $\tilde{w}_h^S \equiv \log w_h^S$. Similarly, let F^I be the distribution of log h-type wages among inclusive firms, i.e. $\tilde{w}_h^I \sim F^I$. Expressing these distributions in terms of log (rather than dollar) wages will simplify the proofs below. It then follows that:

$$F^S(\tilde{w}) = F^x\left(\tilde{w} - \tilde{w}_h^S(1)\right) \tag{C6}$$

$$F^I(\tilde{w}) = F^x\left(\tilde{w} - \tilde{w}_h^I(1)\right) \tag{C7}$$

where $\tilde{w}_h^S(1) = \log \frac{\varepsilon}{1+\varepsilon} p_h$ and $\tilde{w}_h^I(1) = \log \frac{1+\frac{1}{\beta}\cdot\phi^{1+\varepsilon}\frac{\Omega_l}{\Omega_h}}{1+\phi^{1+\varepsilon}\frac{\Omega_l}{\Omega_h}} \frac{\varepsilon}{1+\varepsilon} p_h$. In summary, both the F^S and F^I distributions have identical variance (equal to ν, the same as for firm productivity x), but inclusive firms have a lower mean.

We now turn to the labor supply intercepts, Ω_h and Ω_l. Using equation (3), the intercept for h-type workers can be written as:

$$\begin{aligned}\Omega_h &= \frac{n_h}{k}\left[\sigma\int_{\tilde{w}} e^{\varepsilon\tilde{w}} dF^S(\tilde{w}) + (1-\sigma)\int_{\tilde{w}} e^{\varepsilon\tilde{w}} dF^I(\tilde{w})\right]^{-1} \\ &= \frac{n_h}{k}\left[\sigma\left(w_h^S(1)\right)^{\varepsilon} + (1-\sigma)\left(w_h^I(1)\right)^{\varepsilon}\right]^{-1}\left[\int_{\tilde{x}} e^{\varepsilon\tilde{x}} dF^x(\tilde{x})\right]^{-1}\end{aligned} \tag{C8}$$

where n_h is the measure of h-type workers, and k is the measure of firms. In the first line of (C8), the square brackets contain an average of the wages (with an ε exponent) of selective firms (weighted by the selective firm share σ) and inclusive firms (weighted $1-\sigma$). The second line follows from the definition of (C6) and (C7), as well as the fact that $\tilde{w}_h^S(x) = \tilde{w}_h^S(1) + \tilde{x}$ and $\tilde{w}_h^I(x) = \tilde{w}_h^I(1) + \tilde{x}$: this additive separability allows us to disentangle the x terms from the rest of the expression. Similarly, the labor supply intercept for l-type workers can be written as:

$$\begin{aligned}\Omega_l &= \frac{n_l}{k}\left[(1-\sigma)\int_{\tilde{w}} \phi^{\varepsilon} e^{\varepsilon\tilde{w}} dF^I(\tilde{w})\right]^{-1} \\ &= \frac{n_l}{k}\left[(1-\sigma)\left(w_h^I(1)\right)^{\varepsilon}\right]^{-1}\left[\int_{\tilde{x}} e^{\varepsilon\tilde{x}} dF^x(\tilde{x})\right]^{-1}\end{aligned} \tag{C9}$$

Putting these together, the intercept ratio is identical to equation (B18) in the baseline

model:

$$\frac{\Omega_l}{\Omega_h} = \frac{1-\alpha}{\alpha} \cdot \frac{\beta}{\phi^{1+\varepsilon}} \left[1 + \frac{\sigma}{1-\sigma} \left(\frac{w_h^I(1)}{w_h^I(1)} \right)^{\varepsilon} \right] = \frac{1-\alpha}{\alpha-\sigma} \cdot \frac{\beta}{\phi^{1+\varepsilon}} \tag{C10}$$

And hence, the equilibrium selective share σ will take an identical form to the baseline model, as specified by equation (B19). Finally, inserting (C10) into (C3), the log wage of an inclusive firm with productivity x can be written as:

$$\tilde{w}_h^I(x) = \tilde{w}_h^S(x) - \kappa \tag{C11}$$

where k is the pay differential between equally productive inclusive and selective firms:

$$\kappa = \log \left[1 + \frac{1-\alpha}{1-\sigma} (\beta - 1) \right] \tag{C12}$$

C.2 Proof of Proposition 2b: Firm size effects

Proposition 2b states that log employment is initially positive and concave (and possibly hump-shaped) in log firm pay. The key insight underlying these results is that firm pay may vary for two reasons: (i) heterogeneity in productivity x and (ii) choice of selective or inclusive strategy. As we have shown above, these two sources of variation are orthogonal: a firm's productivity x has no effect on whether it adopts a selective strategy.

It is first useful to define the selective share $\sigma(\tilde{w}_h)$, among firms offering a log h-type wage equal to $\tilde{w}_h$:

$$\begin{aligned} \sigma(\tilde{w}_h) &= \frac{\sigma f^S(\tilde{w}_h)}{(1-\sigma) f^I(\tilde{w}_h) + \sigma f^S(\tilde{w}_h)} \\ &= \frac{\sigma f^S(\tilde{w}_h)}{(1-\sigma) f^S(\tilde{w}_h + \kappa) + \sigma f^S(\tilde{w}_h)} \\ &= \left[\frac{1-\sigma}{\sigma} \exp\left(-\frac{\kappa}{\nu^2} \left(\tilde{w}_h - \tilde{w}_h^S(1) + \frac{1}{2}\kappa \right) \right) + 1 \right]^{-1} \end{aligned} \tag{C13}$$

where σ is the unconditional selective share, and $\tilde{w}_h^S(1) = \log \frac{\varepsilon}{1+\varepsilon} p_h$. The second equality follows from the definitions of F^S and F^I in (C6) and (C7), and the definition of k in equation (C12): i.e., the pay differential between equally productive selective and inclusive firms. The final equality follows from the fact that F^S and F^I are normally distributed, with means $\tilde{w}_h^S(1)$ and $\tilde{w}_h^S(1) - k$ respectively, and variance ν^2. Equation (C13) shows that the selective share $\sigma(\tilde{w}_h)$ is increasing in firm pay, and varies from 0 (for very low $\tilde{w}_h$) to 1 (for very high $\tilde{w}_h$). Intuitively, selective firms pay higher wages (conditional on productivity x); so the higher up the pay distribution we move, the larger the representation of selective firms.

Next, we consider how log firm employment varies over the firm pay distribution. Let $E\left[\log l|\tilde{w}_h\right]$ denote the expectation of log firm employment, conditional on the firm offering a log h-type wage equal to $\tilde{w}_h$. This is a weighted average of the expected log employment of selective and inclusive firms, with weights equal to the selective and inclusive shares at $\tilde{w}_h$:

$$E\left[\log l|\tilde{w}_h\right] = \sigma\left(\tilde{w}_h\right) E\left[\log l^S|\tilde{w}_h\right] + \left[1-\sigma\left(\tilde{w}_h\right)\right] \log\left[\log l^I|\tilde{w}_h\right] \tag{C14}$$

where $\sigma\left(\tilde{w}_h\right)$ is defined by (C13). Since selective firms recruit all h-type workers who are willing to work at $\tilde{w}_h$, their expected employment is simply equal to the h-type labor supply curve. And since inclusive firms recruit all workers (both h- and l-type) who are willing to work at $\tilde{w}_h$, their expected employment is equal to the sum of the h- and l-type labor supply curves. So we have:

$$E\left[\log l|\tilde{w}_h\right] = \sigma\left(\tilde{w}_h\right) \log l_h\left(e^{\varepsilon\tilde{w}_h}\right) + \left[1-\sigma\left(\tilde{w}_h\right)\right] \log\left[l_h\left(e^{\varepsilon\tilde{w}_h}\right) + \log l_l\left(\phi^{\varepsilon} e^{\varepsilon\tilde{w}_h}\right)\right] \tag{C15}$$

Inserting the labor supply curve (2) and rearranging:

$$E\left[\log l|\tilde{w}_h\right] = \log\left(\Omega_h + \Omega_l\phi^{\varepsilon}\right) + \varepsilon\tilde{w}_h - \log\left(1+\frac{\Omega_l\phi^{\varepsilon}}{\Omega_h}\right)\sigma\left(\tilde{w}_h\right) \tag{C16}$$

The first term on the right-hand side is a constant. The second term is linear and increasing in $\tilde{w}_h$, with slope ε: this is the contribution of the upward-sloping labor supply curve (high-paying firms attract more workers). The final term is decreasing in the selective share $\sigma\left(\tilde{w}_h\right)$: at higher firm pay $\tilde{w}_h$, a larger share of firms are selective, so there is more rationing of l-types.

The first derivative of $E\left[\log l|\tilde{w}_h\right]$ can be written as:

$$\frac{d}{d\tilde{w}_h} E\left[\log l|\tilde{w}_h\right] = \varepsilon - \frac{k}{\nu^2}\log\left(1+\frac{\Omega_l\phi^{\varepsilon}}{\Omega_h}\right)\sigma\left(\tilde{w}_h\right)\left[1-\sigma\left(\tilde{w}_h\right)\right] \tag{C17}$$

As $\tilde{w}_h$ becomes small, the selective share $\sigma\left(\tilde{w}_h\right)$ goes to zero, and the derivative converges to the labor supply elasticity ε. But for larger $\tilde{w}_h$, the second term ensures that the derivative drops below ε. The second derivative can be written as:

$$\frac{d^2}{d\tilde{w}_h^2} E\left[\log l|\tilde{w}_h\right] = -\frac{k}{\nu^2}\log\left(1+\frac{\Omega_l\phi^{\varepsilon}}{\Omega_h}\right)\sigma\left(\tilde{w}_h\right)\left[1-\sigma\left(\tilde{w}_h\right)\right]\left[1-2\sigma\left(\tilde{w}_h\right)\right] \tag{C18}$$

which is negative for sufficiently small $\tilde{w}_h$. This proves Proposition 2b: log employment is initially positive and concave (and possibly hump-shaped) in log firm pay.

Finally, notice the curvature of $E\left[\log l|\tilde{w}_h\right]$ is more substantial (and more likely to be

hump-shaped) if the ratio $\frac{\kappa}{\nu^2}$ is larger. Recall that κ is the pay differential between inclusive and selective firms (for given firm productivity x), and ν^2 is the variance of log firm productivity. A larger κ indicates that the "quality motive" is more dominant for firms choosing high pay (firms seeking more h-type employment), and a larger ν indicates that the "quantity motive" is more important (firms seeking more workers of any type). Intuitively, the more important the quality motive, the stronger the quantity-quality trade-off, and the more substantial the curvature of $E\left[\log l|\tilde{w}_h\right]$.

D Extension with N skill types

D.1 Description of framework

In this appendix, we generalize the baseline model from two to N skill types. Firms choose wages and employment, for every skill type s, to maximize profit:

$$\max_{\{w_s;l_s\}_{s=1}^N} \pi\left(w_1,...,w_N;l_1,...,l_N\right)=\sum_{s=1}^{N}\left(p_s-w_s\right)l_s \tag{C19}$$

where skill s productivity p_s is increasing in s, and skill types are perfect substitutes. Firms are subject to labor supply constraints:

$$l_s \leq l_s\left(w_s\right) \tag{C20}$$

where the labor supply curves $l_s\left(w_s\right)$ are defined by (2), and to pay equity constraints:

$$w_s \geq \phi_s w_N \tag{C21}$$

for every skill type s. The Nth equation of (7) is of course redundant, but it is useful for notation to normalize ϕ_N to 1. Analogous to the baseline model, we can also define the "bite" β_s of each equity constraint as:

$$\beta_s=\phi_s\frac{p_N}{p_s} \tag{C22}$$

where $\beta_N=1$. We assume that β_s is strictly decreasing in s, so the equity constraints bind for all skill types s (since $\beta_s>1$ for $s<N$), and the bite is stronger for less productive workers. This is necessarily the case if there is perfect pay equity ($\phi_s=1$ for all s), or more generally if wages are compressed (within firms) relative to productivity differentials.

D.2 Equilibrium strategies

As in the baseline model, since the equity constraints bind, wages will take log additive form:

$$\log w_{sf} = \eta_f + \lambda_s \tag{C23}$$

where firms choose a common firm effect η_f (equal to w_{Nf} in the model, for the top skill type), and the skill effect $\lambda_s = \log \phi_s$ represents the fixed internal pay differential (which firms take as given).

Consider a firm which offers N-type workers a wage of w_N (which determines the common firm effect). Given the equity constraint, the profit from employing an s-type worker is then equal to:

$$p_s - \phi_s w_N = \left(\frac{1}{\beta_s} - \frac{w_N}{p_N} \right) \phi_s p_N \tag{C24}$$

using equation (C22). Firms will employ all willing s-type workers if $p_s \geq \phi_s w_N$ (so the s-type labor supply constraint will bind), and will employ none if $p_s < \phi_s w_N$. But since the constraint bite β_s is decreasing in s (by assumption), equation (C24) implies that if a firm employs s-type workers, it must also employ all workers with skill exceeding s.

It follows that there are N possible strategies in equilibrium (one corresponding to each skill type), which we index z. Firms adopting strategy z employ all workers with skill $s \geq z$, and reject all workers with skill $s < z$. More formally, let w_s^z denote the optimal wage paid by strategy-z firms to s-type workers, and let l_s^z denote the optimal employment of z-type workers. The labor supply constraints bind, i.e. $l_s^z = l_s\left(w_s^z\right)$, for all skill types $s \geq z$. And optimal employment $l_s^z = 0$ for all skill types $s < z$. Strategy z is internally consistent if hiring workers with skill $s < z$ is unprofitable at the chosen wage, i.e., if the s-type wage $w_s^z = \phi_s w_N^z$ (as fixed by the equity constraint) exceeds their productivity p_s.

As in the baseline model, though firms in the baseline model are identical, they may choose different pay strategies in equilibrium. Let σ^k denote the equilibrium share of firms which choose strategy z. Since all firms must choose one of these N strategies, these shares must sum to 1:

$$\sum_z \sigma^z = 1 \tag{C25}$$

D.3 Optimal wage of strategy-z firm

Strategy-z firms do not employ workers with skill $s < z$, so they are not subject to the equity constraint for these workers. But the labor supply constraints will bind for all skill types

$s \geq z$. We can then re-write the firm's problem in (C19) as:

$$\max_{w_N} \pi^z (w_N) = \sum_{s \geq z}^{N} (p_s - w_s) l_s \tag{C26}$$

The first-order constraint is then:

$$\sum_{s \geq z} \phi_s (p_s - \phi_s w_N) l_s' (\phi_s w_N) = \sum_{s \geq z} \phi_s l_s (\phi_s w_N) \tag{C27}$$

Using the labor supply constraint (2), this implies:

$$w_N^z = \frac{\sum_{s \geq z} \frac{\phi_s}{\beta_s} \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}}{\sum_{s \geq z} \phi_s \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}} \cdot \frac{\varepsilon}{1+\varepsilon} p_N \tag{C28}$$

Finally, using (C26), optimal profit of strategy-k firms is:

$$\pi^z = \frac{\varepsilon^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \cdot \frac{\left(\sum_{s \geq z} \frac{\phi_s}{\beta_s} \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N} \right)^{1+\varepsilon}}{\left(\sum_{s \geq z} \phi_s \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N} \right)^{\varepsilon}} \cdot \Omega_N p_N^{1+\varepsilon} \tag{C29}$$

D.4 Labor supply intercepts

To solve for equilibrium, we next require expressions for the labor supply intercepts Ω_s. Since s-type workers are only employed by firms with strategy $z \leq s$, equation (3) implies:

$$\Omega_s = \frac{n_s}{k} \left[\sum_{z \leq s} \sigma^z (\phi_z w_N^z)^\varepsilon \right]^{-1} \tag{C30}$$

Taking the ratio relative to the top skill types ($S = N$), and weighting by ϕ_s^ε, we have:

$$\frac{\phi_s^\varepsilon \Omega_s}{\Omega_N} = \frac{\alpha_s}{\alpha_N} \cdot \frac{\beta_s}{\phi_s} \cdot \frac{\sum_z \sigma^z \left(\frac{w_N^z}{w_N^N} \right)^\varepsilon}{\sum_{z \leq s} \sigma^z \left(\frac{w_N^z}{w_N^N} \right)^\varepsilon} \tag{C31}$$

where

$$\alpha_s = \frac{\alpha_s p_s}{\sum_x \alpha_x p_x} \tag{C32}$$

is the output share of s-type workers.

Also, from (C28), notice the optimal wage of strategy-N firms is:

$$w_N^N = \frac{\varepsilon}{1+\varepsilon} p_N \tag{C33}$$

So the relative wage $\frac{w_N^z}{w_N^N}$ in equation (C31) is equal to:

$$\frac{w_N^z}{w_N^N} = \frac{\sum_{s \geq z} \frac{\phi_s}{\beta_s} \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}}{\sum_{s \geq z} \phi_s \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}} \tag{C34}$$

for strategy $z < N$.

D.5 Equilibrium

In equilibrium, as long as $\alpha_1 > 0$, at least some firms must opt for strategy 1: i.e., $\sigma^1 > 0$. This is because type-1 workers are only employed by strategy-1 firms, and these workers cannot be left unemployed in equilibrium. Otherwise, the profit from strategy 1 would exceed all others, so at least some firms must adopt this strategy (a contradiction).

For all other strategies z, there are two possibilities. Either no firms adopt strategy z, so we have:

$$\sigma^z = 0 \tag{C35}$$

This requires that strategy z is less profitable than strategy 1 (i.e. $\pi^z < \pi^1$). Or alternatively, at least some firms adopt strategy z (i.e. $\sigma^z > 0$), which requires that strategies z and 1 are equally profitable (i.e. $\pi^z = \pi^1$). From equation (C29), equal profits implies:

$$\frac{\left(\sum_{s \geq z} \frac{\phi_s}{\beta_s} \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}\right)^{1+\varepsilon}}{\left(\sum_{s \geq z} \phi_s \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}\right)^{\varepsilon}} = \frac{\left(\sum_{s \geq z} \frac{\phi_s}{\beta_s} \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}\right)^{1+\varepsilon}}{\left(\sum_{s \geq z} \phi_s \cdot \frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}\right)^{\varepsilon}} \tag{C36}$$

In equilibrium, we then have $3N-2$ unknowns: (i) the strategy shares σ^z for $z = 1, ..., N$; (ii) the optimal wages $\frac{w_N^z}{w_N^N}$ for strategies $z = 1, ..., N-1$ (relative to the strategy-N wage); and (iii) the relative labor supply intercepts $\frac{\phi_s^\varepsilon \Omega_s}{\Omega_N}$ for skill types $s = 1, ..., N-1$. And we also have $3N-2$ equations: (i) the relative intercept equations (C31) for strategies $z = 1, ..., N-1$; (ii) the relative wage equations (C34) for strategies $z = 1, ..., N-1$; (iii) equation (C25), which ensures the strategy shares sum to 1; and finally, (iv) we have one equilibrium condition for every strategy $z = 2, ..., N$: either (C35) or (C36).

E Quantification of model's parameters

This appendix provides the technical details for quantifying the model parameters in Section 4.3. We implement this exercise in an extension with heterogeneous firms (as in Appendix C) and three skill types (a special case of Appendix D).

The three skill types correspond to non-graduates, non-STEM graduates and STEM graduates, and we denote them l, m and h, respectively. Assuming the equity constraint has stronger bite for l-types, i.e., $\beta_l > \beta_m$ (we will validate this assumption ex post), Appendix D shows that firms may pursue one of three strategies in equilibrium: (i) L-strategy: hire all willing workers, and pay wages w_s^L to s-type workers; (ii) M-strategy: hire only m- and h-type workers, and pay wages w_s^M to s-type workers; and (iii) H-strategy: hire only h-type workers, and pay them w_h^H. Let σ^L, σ^M and σ^H denote the equilibrium shares of L, M and H-strategy firms, where $\sigma^L + \sigma^M + \sigma^H = 1$.

As in Appendix C, the marginal product of s-type workers in firm f is equal to $p_{sf} = x_f p_s$, where $\log x_f$ is distributed normally across firms with mean 0 and variance ν.

E.1 Solution method: Step 1

To solve for the parameter values, we iterate over two steps. In the first step, for given labor supply elasticity ε, we solve for six parameters, using six moments and six equations. The six parameters are: $\frac{w_h^L}{w_h^H}$, $\frac{w_h^M}{w_h^H}$, $\frac{\phi_l^\varepsilon \Omega_l}{\Omega_h}$, $\frac{\phi_m^\varepsilon \Omega_m}{\Omega_h}$, σ^L, σ^M; and the six moments are: ϕ_m, ϕ_l, $\frac{n_m}{n_h}$,$\frac{n_l}{n_h}$, $E[\log w_m] - E[\log w_h]$, $E[\log w_l] - E[\log w_h]$.

We now set out the six equations. Recall from Appendix C that optimal wages (by strategy) and profits are log additive in firm productivity. It follows that the intercept ratios, i.e. $\frac{\Omega_l}{\Omega_h}$ and $\frac{\Omega_m}{\Omega_h}$, are independent of the firm productivity distribution; and the equilibrium strategy shares (i.e. σ^L, σ^M and σ^H) are orthogonal to firm productivity. We can therefore solve for $\frac{w_h^L}{w_h^H}$, $\frac{w_h^M}{w_h^H}$, $\frac{\phi_l^\varepsilon \Omega_l}{\Omega_h}$, $\frac{\phi_m^\varepsilon \Omega_m}{\Omega_h}$, σ^L and σ^M independently of the firm productivity distribution.

We have two equilibrium conditions for equal profits, which follow from equation (C36) in the N-type model. Equal profits for the L- and H-strategies implies:

$$\frac{w_h^L}{w_h^H} = \left(1 + \phi_m \frac{\phi_m^\varepsilon \Omega_m}{\Omega_h} + \phi_l \frac{\phi_l^\varepsilon \Omega_l}{\Omega_h}\right)^{-\frac{1}{1+\varepsilon}} \tag{D1}$$

and equal profits for the M- and H-strategies implies:

$$\frac{w_h^M}{w_h^H} = \left(1 + \phi_m \cdot \frac{\phi_m^\varepsilon \Omega_m}{\Omega_h}\right)^{-\frac{1}{1+\varepsilon}} \tag{D2}$$

Next, we have two equations for equilibrium ratios of the labor supply intercepts. From equation (C31) in the N-type model, these are:

$$\frac{\phi_l^{\varepsilon}\Omega_l}{\Omega_h} = \frac{n_l}{n_h} \cdot \frac{1 + \frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + \frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}}{\frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}} \tag{D3}$$

and

$$\frac{\phi_m^{\varepsilon}\Omega_m}{\Omega_h} = \frac{n_m}{n_h} \cdot \frac{1 + \frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + \frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}}{\frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + \frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}} \tag{D4}$$

Finally, we use two expressions for the expected log wages of l-types and m-types, expressed relative to h-types. These are:

$$E\left[\log w_l\right] - E\left[\log w_h\right] = \log \phi_l + \frac{\frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^L}{w_h^M} + \log \frac{w_h^L}{w_h^H}}{\frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + 1} \tag{D5}$$

and

$$\begin{aligned} E\left[\log w_m\right] - E\left[\log w_h\right] \;=\; & \log \phi_m + \frac{\frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^L}{w_h^H} + \frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^M}{w_h^H}}{\frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon}} \\ & - \frac{\frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^L}{w_h^H} + \frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^M}{w_h^H}}{\frac{\sigma_l}{\sigma_h}\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \frac{\sigma_m}{\sigma_h}\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + 1} \end{aligned} \tag{D6}$$

E.2 Solution method: Step 2

In the second step, we pick the firm productivity variance ν and the labor supply elasticity ε to match two additional moments: (i) the elasticity of firm size with respect to AKM firm effects at the bottom quartile of the firm pay distribution (denoted ε_{bottom}) and (ii) the variance of AKM firm effects (var_{AKM}). To estimate these moments in the model, we first simulate a panel of 1 million firms, drawing log firm productivity $\tilde{x}_f$ from a normal distribution with mean zero and variance ν.

For each simulated firm f, we compute employment and wages by education group, organize this data in "long" form (with each row corresponding to a firm $\times$ education group), and then estimate an AKM model by regressing log wages on firm and education fixed effects, and save the firm premia as η_f. We then regress log employment on the firm effects η_f across

firms in the bottom quartile of the η_f distribution:

$$\log l_f = \alpha + \gamma\eta_f + \epsilon_f, \quad \text{for rank}(\eta_f) \leq 0.25 \tag{D7}$$

The estimated coefficient γ provides our model-based moment for ε_{bottom}. The variance of the estimated firm effects η_f across all firms provides our model-based moment for var_{AKM}.

Following San (2023), we implement an iterative gradient descent procedure to find values of the firm productivity variance ν and labor supply elasticity ε that equate the model-based and empirical moments. The procedure updates parameters in each iteration according to the moments most affected by those parameters, based on the model's structure. Specifically, at each iteration i, we:

1. Compute model moments $m_i = (m_{i1}, m_{i2})$ for current parameter values $\theta_i = (\varepsilon_i, \nu_i)$.
2. Update parameters according to $\theta_{i+1} = \theta_i + \eta(m^* - m_i)$, where $m^* = (\varepsilon_{bottom}, var_{AKM})$ are the empirical target moments and η is the learning rate

The algorithm continues until the distance between model and empirical moments falls below a tolerance level τ: i.e. $\sum_j |m_{ij} - m_j^*| < \tau$. We set the learning rate $\eta = 0.1$ and tolerance $\tau = 10^{-3}$. At each iteration, we resolve the equilibrium equations from Step 1 given the updated ε.

The final estimated parameters imply a labor supply elasticity of $\varepsilon = 3.78$ and productivity variance of $\nu = 0.02$. With these values, the model successfully replicates both the firm size-wage premium relationship at the bottom of the firm distribution ($\varepsilon_{bottom} = 3.61$ in both model and data) and the overall dispersion in firm wage premia ($var_{AKM} = 0.0355$ in both model and data).

F Derivation of counterfactual outcomes

In this appendix, we derive expressions for the impact of two counterfactuals in a model with three skill types $s = \{l, m, h\}$. We consider the removal of the equity constraint in Appendix F.1 and the prohibition of selective hiring strategies in Appendix F.2.

F.1 Counterfactual with no equity constraint

Impact on expected log wages

In the counterfactual, all workers earn the unconstrained optimum wage, for $s = \{l, m, h\}$, in all firms. Denoting counterfactual outcomes with a $CF1$ superscript, wages for skill type

s are therefore:

$$w_s^{CF1} = w_h^* = \frac{\varepsilon}{1+\varepsilon} p_s \tag{D8}$$

We now derive the impact on expected log wages for each skill type. Since h-types are employed by all firms in the baseline model, the counterfactual impact can be written as:

$$\begin{aligned} \log w_h^{CF1} - E\left[\log w_h\right] &= \log w_h^{CF} \\ &\quad - \frac{\sigma^L l_h\left(w_h^L\right)\log w_h^L + \sigma^M l_h\left(w_h^M\right)\log w_h^M + \sigma^H l_h\left(w_h^H\right)\log w_h^H}{\sigma^L l_h\left(w_h^L\right) + \sigma^M l_h\left(w_h^M\right) + \sigma^H l_h\left(w_h^H\right)} \\ &= - \frac{\sigma^L \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^L}{w_h^H} + \sigma^M \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^M}{w_h^H}}{\sigma^L \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \sigma^M \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + \sigma^H} \end{aligned} \tag{D9}$$

where σ^L, σ^M and σ^H are the shares of L, M and H-strategy firms respectively (using the notation of Appendix D). The second line uses the labor supply function in (2), and also the fact the H-strategy wage w_h^H is equal to the unconstrained optimum w_h^* (since these firms hire only h-types).

The impact on m-types is:

$$\begin{aligned} \log w_m^{CF1} - E\left[\log w_m\right] &= \left(\log w_m^{CF1} - \log w_h^{CF1}\right) + \left(\log w_h^{CF1} - E\left[\log w_h\right]\right) \\ &\quad + \left(E\left[\log w_h\right] - E\left[\log w_m\right]\right) \\ &= - \log \beta_m - \frac{\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^L}{w_h^H} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^M}{w_h^H}}{\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon}} \end{aligned} \tag{D10}$$

where the second line uses the definition of β_s in (C22), the expected skill differential in (D6), and the counterfactual impact in (D9).

Finally, the impact on l-types is:

$$\begin{aligned} \log w_l^{CF1} - E\left[\log w_l\right] &= \left(\log w_l^{CF1} - \log w_h^{CF1}\right) + \left(\log w_h^{CF1} - E\left[\log w_h\right]\right) \\ &\quad + \left(E\left[\log w_h\right] - E\left[\log w_l\right]\right) \\ &= - \log \beta_l - \log \frac{w_h^L}{w_h^H} \end{aligned} \tag{D11}$$

where the second line uses the expected skill differential in (D5) and the counterfactual impact in (D9).

Impact on expected utility

We now turn to the effect on expected utility. Note we weight utility by $\frac{1}{\varepsilon}$ for this exercise, to ensure it is in log wage units: see equation (1). Using equation (B27), the impact on h-type utility can be written as:

$$\begin{aligned}\frac{1}{\varepsilon}\left(\bar{u}_h^{CF1}-\bar{u}_h\right) &= \frac{1}{\varepsilon}\log\frac{\left(w_h^{CF1}\right)^{\varepsilon}}{\sigma_l\left(w_h^L\right)^{\varepsilon}+\sigma_m\left(w_h^M\right)^{\varepsilon}+\sigma_h\left(w_h^H\right)^{\varepsilon}} \\ &= -\frac{1}{\varepsilon}\log\left[\sigma_l\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}+\sigma_m\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon}+\sigma_h\right]\end{aligned} \tag{D12}$$

which again uses the equality between w_h^{CF1} and w_h^H. For m-types, the impact is:

$$\begin{aligned}\frac{1}{\varepsilon}\left(\bar{u}_m^{CF1}-\bar{u}_m\right) &= \frac{1}{\varepsilon}\log\frac{\left(w_m^{CF1}\right)^{\varepsilon}}{\sigma_l\phi_m^{\varepsilon}\left(w_h^L\right)^{\varepsilon}+\sigma_m\phi_m^{\varepsilon}\left(w_h^M\right)^{\varepsilon}} \\ &= -\log\beta_m-\frac{1}{\varepsilon}\log\left[\sigma_l\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}+\sigma_m\left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon}\right]\end{aligned} \tag{D13}$$

And for l-types:

$$\frac{1}{\varepsilon}\left(\bar{u}_l^{CF1}-\bar{u}_l\right) = -\log\beta_l-\frac{1}{\varepsilon}\log\left[\sigma_l\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}\right] \tag{D14}$$

The impact on expected amenities (weighted by $\frac{1}{\varepsilon}$) is simply the difference between the expected utility and log wage effects.

F.2 Counterfactual with no selective strategy

Effects on expected log wages

In this counterfactual, all firms adopt the inclusive L-strategy, and employ all workers who are willing to work: i.e., the labor supply constraints always bind. Building from the N-types case in equation (C34), the optimal L-strategy wage (for h-type workers) can be written as:

$$w_h^{CF2}=w_h^L=\frac{1+\frac{\phi_m}{\beta_m}\cdot\frac{\phi_m^{\varepsilon}\Omega_m}{\Omega_h}+\frac{\phi_l}{\beta_l}\cdot\frac{\phi_l^{\varepsilon}\Omega_l}{\Omega_h}}{1+\phi_m\frac{\phi_m^{\varepsilon}\Omega_m}{\Omega_h}+\phi_l\frac{\phi_l^{\varepsilon}\Omega_l}{\Omega_h}}\cdot\frac{\varepsilon}{1+\varepsilon}p_h \tag{D15}$$

Since all firms adopt the same strategy (and pay the same wage), the labor supply intercept ratios collapse to the aggregate employment ratios:

$$\frac{\phi_s^{\varepsilon}\Omega_s}{\Omega_h}=\frac{n_s}{n_h} \tag{D16}$$

for $s = \{l, m\}$. More formally, this can be seen from equation (C30) in the N-type model. Imposing this on the equation above, we have:

$$w_h^{CF2} = \frac{1 + \frac{\phi_m}{\beta_m} \cdot \frac{n_m}{n_h} + \frac{\phi_l}{\beta_l} \cdot \frac{n_l}{n_h}}{1 + \phi_m \cdot \frac{n_m}{n_h} + \phi_l \cdot \frac{n_l}{n_h}} \cdot \frac{\varepsilon}{1+\varepsilon} p_h \tag{D17}$$

Imposing the equity constraints, m-types receive $\phi_m w_h^{CF2}$ and l-types receive $\phi_l w_h^{CF2}$.

Building from (D9), the impact on the expected log h-type wage is:

$$\log w_h^{CF2} - E\left[\log w_h\right] = \log \frac{1 + \frac{\phi_m}{\beta_m} \cdot \frac{n_m}{n_h} + \frac{\phi_l}{\beta_l} \cdot \frac{n_l}{n_h}}{1 + \phi_m \cdot \frac{n_m}{n_h} + \phi_l \cdot \frac{n_l}{n_h}} - \frac{\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^L}{w_h^H} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^M}{w_h^H}}{\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + \sigma_h} \tag{D18}$$

For m-types, we have:

$$\log w_m^{CF2} - E\left[\log w_m\right] = \log \frac{1 + \frac{\phi_m}{\beta_m} \cdot \frac{n_m}{n_h} + \frac{\phi_l}{\beta_l} \cdot \frac{n_l}{n_h}}{1 + \phi_m \cdot \frac{n_m}{n_h} + \phi_l \cdot \frac{n_l}{n_h}} - \frac{\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^L}{w_h^H} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} \log \frac{w_h^M}{w_h^H}}{\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon}} \tag{D19}$$

And for l-types:

$$\log w_l^{CF2} - E\left[\log w_l\right] = \log \frac{1 + \frac{\phi_m}{\beta_m} \cdot \frac{n_m}{n_h} + \frac{\phi_l}{\beta_l} \cdot \frac{n_l}{n_h}}{1 + \phi_m \cdot \frac{n_m}{n_h} + \phi_l \cdot \frac{n_l}{n_h}} - \log \frac{w_h^L}{w_h^H} \tag{D20}$$

Impact on expected utility

We now turn to the effect on expected utility. Note we weight utility by $\frac{1}{\varepsilon}$ for this exercise, to ensure it is in log wage units: see equation (1). Building from (D12), the impact on expected h-type utility is:

$$\frac{1}{\varepsilon}\left(\bar{u}_h^{CF2} - \bar{u}_h\right) = \log \frac{1 + \frac{\phi_m}{\beta_m} \cdot \frac{n_m}{n_h} + \frac{\phi_l}{\beta_l} \cdot \frac{n_l}{n_h}}{1 + \phi_m \cdot \frac{n_m}{n_h} + \phi_l \cdot \frac{n_l}{n_h}} - \frac{1}{\varepsilon} \log \left[\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon} + \sigma_h\right] \tag{D21}$$

For m-types, we have:

$$\frac{1}{\varepsilon}\left(\bar{u}_m^{CF2} - \bar{u}_m\right) = \log \frac{1 + \frac{\phi_m}{\beta_m} \cdot \frac{n_m}{n_h} + \frac{\phi_l}{\beta_l} \cdot \frac{n_l}{n_h}}{1 + \phi_m \cdot \frac{n_m}{n_h} + \phi_l \cdot \frac{n_l}{n_h}} - \frac{1}{\varepsilon} \log \left[\sigma_l \left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon} + \sigma_m \left(\frac{w_h^M}{w_h^H}\right)^{\varepsilon}\right] \tag{D22}$$

And for l-types:

$$\frac{1}{\varepsilon}\left(\bar{u}_l^{CF2} - \bar{u}_l\right) = \log\frac{1+\frac{\phi_m}{\beta_m}\cdot\frac{n_m}{n_h}+\frac{\phi_l}{\beta_l}\cdot\frac{n_l}{n_h}}{1+\phi_m\cdot\frac{n_m}{n_h}+\phi_l\cdot\frac{n_l}{n_h}} - \frac{1}{\varepsilon}\log\left[\sigma_l\left(\frac{w_h^L}{w_h^H}\right)^{\varepsilon}\right] \tag{D23}$$

As before, the impact on expected amenities (weighted by $\frac{1}{\varepsilon}$) is the difference between the expected utility and log wage effects.

G Alternative models

In this appendix, we describe three alternative models that we compare to our baseline equity constraint model. For each model, we explain the key differences and how we calibrate the parameters.

G.1 Model 2: Skill-neutral firm heterogeneity

Our first alternative specification removes the equity constraint and assumes firms differ only in skill-neutral productivity.

Since there is no equity constraint, firms pay the unconstrained optimum to each skill type $s = \{h, m, l\}$, in line with equation (10). For skill s, the optimal wage is:

$$w_{sf} = \frac{\varepsilon}{1+\varepsilon}p_{sf} \tag{D24}$$

where p_{sf} is the marginal product of s-type workers in firm f:

$$p_{sf} = x_f p_s \tag{D25}$$

where x_f is distributed log-normally across firms, with mean 0 and variance ν; and p_s represents base productivity for skill type s.

Labor supply of skill type s to a firm paying wage w is:

$$l_s(w) = \Omega_s w^{\varepsilon} \tag{D26}$$

and the labor supply intercept is:

$$\Omega_s = \left(\sum_f w_{sf}^{\varepsilon}\right)^{-1} n_s \tag{D27}$$

Expressing outcomes relative to h-types, the model can be summarized by six parameters. First, the labor supply elasticity ε determines workers' responsiveness to wage differences across firms. Second, the variance of firm productivity ν governs pay dispersion across firms. The next two are the base productivity differentials, $\log(p_m/p_h)$ and $\log(p_l/p_h)$, which capture skill-specific differences in worker productivity. And finally, the relative labor supply intercepts, Ω_m/Ω_h and Ω_l/Ω_h, reflect the relative availability of each worker type adjusted for their outside options.

We calibrate these parameters to match six empirical moments. The first two are the elasticity of firm size with respect to AKM firm premia at the bottom quartile of the firm pay distribution ($\varepsilon_{bottom} = 3.61$) and the variance of AKM firm effects ($var_{AKM} = 0.0355$): these help identify the labor supply elasticity ε and productivity variance ν. The next two moments are the mean log wage differentials between education groups: between non-STEM and STEM graduates, i.e. $E[\log w_m] - E[\log w_h] = -0.67$, and between non-graduates and STEM graduates, i.e. $E[\log w_l] - E[\log w_h] = -0.47$. The final two moments are the aggregate employment ratios: the ratio of non-STEM to STEM graduate employment, i.e., $n_m/n_h = 3.81$, and the ratio of non-graduate to STEM graduate employment, i.e., $n_l/n_h = 6.85$.

In all three alternative models, we follow a similar two-step estimation procedure to the baseline model. In Step 1, for a given labor supply elasticity ε, we solve for the other parameters to match the wage differentials and aggregate employment ratios. In Step 2, we update ε and ν based on the firm size-wage premia elasticity and AKM variance moments. The calibrated parameters are reported in Table A1.

G.2 Model 3: Skill-biased firm heterogeneity

Our second alternative model allows for skill-biased productivity differences across firms. The marginal product of s-type workers in firm f is:

$$p_{sf} = x_f^{\theta_s} p_s \tag{D28}$$

where the θ_s are skill-specific productivity elasticities with respect to firm heterogeneity x_f. We normalize $\theta_h = 1$ and estimate θ_m and θ_l. The rest of the model is identical to the previous model.

The model can be characterized by eight parameters: the labor supply elasticity ε, the variance of firm productivity ν, the skill-specific productivity elasticities θ_m and θ_l, the base productivity differentials $\log(p_m/p_h)$ and $\log(p_l/p_h)$, and the relative labor supply intercepts, Ω_m/Ω_h and Ω_l/Ω_h.

We calibrate these parameters to match the same eight moments as in the baseline model: the elasticity of firm size with respect to AKM firm effects in the bottom firm quartile (ε_{bottom}), the variance of AKM firm effects (var_{AKM}), the mean firm AKM effects by education group relative to STEM graduates (AKM_{fm} and AKM_{fl}), the log wage differentials between education groups, i.e. $E[\log w_m] - E[\log w_h]$ and $E[\log w_l] - E[\log w_h]$, and the aggregate employment ratios, i.e. n_m/n_h and n_l/n_h.

G.3 Model 4: Skill-varying labor supply elasticities

The final model imposes skill-neutral heterogeneity in firm productivity, but permits the labor supply elasticity to vary by skill group. The utility of worker i of skill type s in firm f now takes the form:

$$u_{isf} = \varepsilon_s \log w_{sf} + a_{if} \tag{D29}$$

where ε_s is the skill-specific labor supply elasticity. Like the previous model, this model also has eight parameters: the base labor supply elasticity ε_h, the elasticity differentials $\varepsilon_m - \varepsilon_h$ and $\varepsilon_l - \varepsilon_h$, the variance of firm productivity ν, the base productivity differentials, $\log(p_m/p_h)$ and $\log(p_l/p_h)$, and the labor supply intercepts, Ω_m and Ω_l.

We calibrate these parameters to match the same eight moments as Model 3 (and the baseline model): the elasticity of firm size with respect to AKM effects (ε_{bottom}), the variance of AKM firm effects (var_{AKM}), the mean firm AKM effects by education group relative to STEM graduates (AKM_{fm} and AKM_{fl}), the log wage differentials between education groups, i.e. $E[\log w_m] - E[\log w_h]$ and $E[\log w_l] - E[\log w_h]$, and the aggregate employment ratios, i.e. n_m/n_h and n_l/n_h.

G.4 Estimation results

For each model, we implement an iterative procedure similar to our baseline model, using gradient descent with step size $\eta = 0.1$ and convergence tolerance of 10^{-3}. The calibrated parameters for all model variants are presented in Table A1. Model 2 yields a labor supply elasticity of 3.61 and sizable productivity gaps across education groups. Model 3 generates substantial skill-biased productivity differences across firms, with larger productivity heterogeneity for high-skilled workers. Model 4 produces considerable heterogeneity in labor supply elasticities across skill groups, with high-skilled workers being the most responsive to wage differences.

H Preparation of Israeli administrative data

This appendix provides additional details on the data preparation and definitions of key variables.

Wages: Our raw data contains observations at the worker × firm × year level, with monthly employment indicators and total annual compensation for each employment spell. We implement several data cleaning procedures to ensure accurate wage measurements: (i) removing observations with missing worker or firm identifiers, (ii) standardizing the treatment of monthly indicators by replacing missing values with zeros, (iii) eliminating exact duplicates across all variables, and (iv) for cases where worker-firm combinations appear multiple times within a year, consolidating by taking the maximum value of monthly indicators and summing the annual earnings.

From this cleaned dataset, we construct an annual panel by assigning individuals to the firm where they worked during November. For each worker-firm match, we impute a monthly salary by dividing total annual earnings by the number of months employed at that firm. In cases where workers had multiple employers in November, we assign the worker to the firm paying the higher monthly salary.

To focus on workers with substantial labor market attachment, we exclude worker-year observations with monthly earnings below 25% of the national average wage that year.[1] Our final sample spans 1990-2019 and includes workers aged 25-64 in each year.

Education: We use the Central Bureau of Statistics (CBS) education registry to classify workers into three mutually exclusive and time-invariant education categories, based on the highest degree they obtained during our sample period. These categories are: (i) non-graduate (no BA-equivalent or higher qualification), (ii) non-STEM graduate (BA-equivalent or higher degree in a non-STEM field), and (iii) STEM graduate (BA-equivalent or higher degree in a STEM field).

Workplace location: We identify workplace locations using multiple data sources. For some analyses, we draw on workplace geographical identifiers from 20% samples of the Israeli census conducted in 1995 and 2008. For years from 2012 onward, we utilize the Arnona (municipal tax) database, which provides detailed location information for businesses.

To assign workers to specific workplace cities (in years from 2012), we implement the following procedure. For each firm, we identify all cities where the firm is registered as paying municipal business tax. We then assign each worker to the workplace location closest to their residence, with residence determined by the city where they pay the highest residential

[1] For context, the statutory minimum wage in Israel ranged between 40-50% of the average wage during our sample period, reaching 48.8% in 2015. Our threshold therefore excludes workers earning approximately half the minimum wage or less, likely representing part-time or marginal employment.

municipality tax. For married individuals who show no tax payments in other locations, we assume they share the same city of residence as their spouse; and for households with no Arnona information, we use their official residence registration according to the Interior Ministry register.

Finally, we aggregate these city-level data into 49 regional units based on Israel's "natural regions" as defined by the Central Bureau of Statistics. These natural regions are constructed to ensure demographic, economic, and social homogeneity of the constituent populations. To ensure sufficient statistical power for all analyses, we merged the three smallest regions into their neighboring regions.

Industry: We use a consistent 2-digit industry classification for each firm across the entire sample period.

I Replication using Veneto Worker History dataset

This appendix reproduces the relationship between log firm size and AKM wage premia using the Veneto Worker History (VWH) dataset, which contains detailed employer-employee linked administrative records for Italy's Veneto region over 1975-2001. The data cover the universe of private sector employment in the region, with particularly good coverage of small establishments that characterize the region's manufacturing sector. We estimate AKM firm effects using log annual earnings for years between 1992 and 2001, and implement the same split-sample approach as in our main analysis to address measurement error.

We plot the relationship in Figure A1, across 20 firm bins ranked by their AKM wage premia. As in the Israeli data, we again see a hump-shaped relationship between firm size and firm wage premia, with employment initially increasing and then decreasing with the wage premium—both for the aggregate data and after residualizing by industry. This suggests that the quantity-quality trade-off is a more general phenomenon, arising from fundamental constraints on firms' wage-setting, rather than from country-specific institutions or policies.

References

San, Shmuel, "Who Works Where and Why: The Role of Social Connections in the Labor Market," 2023. Unpublished.