

Perks Under Pressure? Demand Shocks and Firm Provision of Workplace Amenities*

Manudeep Bhuller[†] Lukas Delgado-Prieto[‡] Samuel Škoda[§]

This Version: October 2026

Abstract

Do firms adjust workplace amenities in response to demand shocks? We develop a monopsony framework in which firms jointly determine wages and amenities, and show that the pass-throughs of demand shocks depend crucially on the structure of amenity provision costs. Using Norwegian administrative data that combine matched employer-employee records, firms' cost accounts, and non-pecuniary attributes from job postings, we provide evidence on how firms adjust amenities in response to demand shocks. Exploiting three alternative sources of variation—abrupt value-added changes, public procurement contract wins, and export demand shifts—we document multi-margin adjustments in non-wage compensation, besides wage and employment. We discuss the implications of our findings for the measurement of employer market power.

Keywords: Demand Shocks, Firm Heterogeneity, Monopsony Power, Workplace Amenities, Non-Pecuniary Job Attributes.

JEL: D22, J31, J42, J51.

* *Acknowledgments:* This project has received funding from the European Union's Horizon Europe Research and Innovation Programme under Grant Agreement No. 101043127. We thank participants at the 2026 CEPR-Banca d'Italia Labor Workshop, the 2026 ELSE Workshop at Bocconi University, the Berlin Applied Micro Seminar, the Catholic University of Milan and the Frisch Centre for helpful comments.

[†]University of Oslo, Statistics Norway, CEPR, IZA, CESifo. E-mail: manudeep.bhuller@econ.uio.no.

[‡]University of Oslo, RFBerlin. E-mail: laprieto@econ.uio.no.

[§]University of Oslo. E-mail: samuel.skoda@econ.uio.no.

1 Introduction

Firms attract and retain workers through wages and non-wage workplace amenities, such as health insurance, flexible scheduling, on-the-job training, and office space, which jointly determine job values. A large empirical literature estimates the pass-through of firm-level demand shocks to wages (Card et al., 2018; Kline et al., 2019; Lamadon et al., 2022), but the adjustments in workplace amenities in response to such shocks remain largely unexplored. This omission matters because, if amenities adjust endogenously alongside wages, both the measurement of employer market power and the welfare interpretation of wage differentials may change—see, e.g., discussions in a recent survey by Caldwell et al. (forthcoming). If part of the labor supply response to firm-level demand shocks operates through adjustments in workplace amenities, then conventional approaches aimed at quantifying wage markdowns may not accurately characterize employer market power. At the same time, whether or not growing firms adjust amenities alongside wages has direct welfare implications.

The endogenous provision of workplace amenities raises three questions that motivate our analysis. First, how do firms decide the optimal mix of wages and amenities in the presence of employer market power? The classical Rosen (1986) compensating differentials model maintains that wages compensate for amenities, implying a negative relationship between the two in a competitive labor market. This implication may not hold when employers have market power, resulting in wage markdowns as well as the under-provision of amenities. Second, what determines the adjustments in wages and amenities in response to firm-level demand shocks? Growing firms have an incentive to adjust both wages and amenities to attract workers, but such adjustments may not be proportional, creating a wedge between the shock pass-throughs in wages and in total compensation. Third, what do conventional estimates of the firm-level labor supply elasticity capture when amenities adjust alongside wages, and how can we quantify welfare-relevant measures of employer market power?

To answer these questions, we develop a monopsony framework in which firms not only have wage-setting power but also determine the provision of workplace amenities. Higher compensation—higher wages or better amenities—attracts more workers, so each firm faces an upward-sloping labor supply curve, as in a classical monopsony (Manning, 2003, 2021). Meanwhile, the mix of wages and amenities offered in the compensation bundle is determined by workers’ valuation of amenities and the amenity provision costs facing the firm. Amenities are multi-dimensional and may differ in their cost structure. For instance, some amenities may have provision costs that are proportional to firm size (e.g., pension premiums), others may involve fixed costs (e.g., on-site facilities), and yet others exhibit congestion or coordination costs (e.g., schedule flexibility). This heterogeneity in amenity cost structure

yields rich comparative statics on how firms respond to shocks. Crucially, firm-level demand shocks affect amenity provision through two channels in our framework: (i) a taste-cost ratio that pins amenity adjustment to the wage pass-through, and (ii) a scale adjustment that alters the amenity provision cost. While amenities with constant per-worker costs track wages, amenities with economies of scale adjust more than wages, and amenities that become costlier to provide at scale adjust less, or even offset wages, if provision costs are sufficiently convex. When amenity provision costs depend on firm size, per-worker *amenity spending* may not track the *amenity level* the employer provides, leading to distinct predictions in our framework about amenity spending and provision. The size-dependent structure of amenity provision costs shapes the distribution of employer market power across firms. In our framework, the ratio of employment-to-wage responses to a firm-level demand shock provides the *effective* labor supply elasticity, incorporating the endogenous amenity responses. The welfare-relevant measure of employer market power is the *total compensation markdown*—the ratio of total compensation to the marginal revenue product—rather than the conventional wage markdown, and the gap between the two is governed by an *amenity cost multiplier*, capturing how total compensation costs scale relative to the wage bill.

Building on this framework, we derive sharp testable predictions and evaluate these using a unique combination of Norwegian data sources. In our empirical evaluation, we must circumvent several methodological challenges. First, amenities are multi-dimensional and hard to measure: most data sources capture only a single margin of adjustment, yet one may expect firms to re-optimize across all margins simultaneously (Mas, 2025). We address this using Norwegian administrative data with information on non-wage components, complemented with the near-universe of job advertisements. This provides a linked dataset with comprehensive amenity coverage, allowing us to test the model’s multi-margin predictions: our data span in-kind compensation, pension premiums, employment conditions (contract types, working-time arrangements), and the amenity content of job advertisements, classified through the text-analysis approach of Audoly et al. (2024). Second, identification requires exogenous variation in product demand at the firm level. We exploit three alternative sources of firm-level demand shifters, relying on the richness of our data and institutional setting: (i) an *internal* shock design that uses abrupt within-firm changes in value added, as in Lamadon et al. (2022), (ii) a *procurement* shock that uses firms winning their first-ever public procurement contract, in the spirit of Kroft et al. (2025) and Volpe (2024), and (iii) a *shift-share export* shock that combines firm-level pre-shock exposure to country-product export markets with shifts in international import demand around the 2007–2009 Great Recession, following Garin and Silvério (2024) and Bhuller et al. (2025). Exploiting these sources of firm-level demand variation allows us to validate our empirical findings across alternative designs.

Third, identification of the total compensation markdowns requires embedding the amenity responses in the structural model. We provide a structural identification and estimation approach that lets us recover the degree of substitutability between wages and amenities in workers' preferences *and* the dispersion of idiosyncratic tastes, which is the underlying source of employer market power in our model. Equipped with amenity cost multipliers and effective labor supply elasticities measured across the firm distribution, we are able to identify the distribution of total compensation markdowns.

Our empirical evidence shows that firm-level demand shocks generate sizable responses across the multiple margins of non-wage compensation that we can measure. Our estimates imply effective labor supply elasticities ranging from roughly 2 to 8 across the distribution of firms, with employment responding more than wages. Firms' total spending on non-wage compensation moves in the same direction as wage costs, with pension premiums and other non-wage personnel costs rising by magnitudes comparable to the wage bill. Meanwhile, per-worker responses differ across cost components. Firms expand coverage of in-kind benefits and raise benefit amounts, shift toward full-time contracts, and tilt job postings toward on-the-job training and workplace quality, while advertised schedule flexibility declines.

Embedding the estimated responses to firm-level demand shocks in our structural model, we recover an elasticity of substitution between wages and amenities in worker preferences that is statistically greater than one but economically close to it, and a structural labor supply elasticity of 4.3, suggesting comparatively modest dispersion of idiosyncratic tastes. Total compensation markdowns rise from roughly 64 to 88 percent of marginal product from the bottom to the top decile of the firm distribution, with wage markdowns another 5.6 to 10.6 percentage points smaller. Our evidence thus suggests that a meaningful share of what might be attributed to employer rent extraction is, in fact, compensation paid through non-wage benefits. We further decompose total compensation markdowns into a *cost-curvature wedge*, reflecting the structure of amenity provision costs, and a *conduct wedge*, reflecting employer market power. Among the least productive firms, about two-thirds of the markdown gap reflects cost curvature, while at the most productive firms conduct accounts for all of it. The conduct wedge itself is stable, with employer market power lowering compensation by 15 to 21 percent below what cost curvature alone would imply. Attributing the entire gap to employer market power would overstate the role of conduct where markdowns are deepest.

Related literature and contribution. Our paper makes four key contributions. First, we develop a tractable monopsony framework in which firms jointly set wages and multi-dimensional amenities facing heterogeneous provision costs that depend on firm size. Second, we show that standard IV estimands using firm-level demand shocks identify the distribution

of effective labor supply elasticities—which incorporate the endogenous amenity responses—rather than the structural dispersion of idiosyncratic tastes, and provide conditions under which the distribution of welfare-relevant total compensation markdowns can be identified from the data. Third, we provide direct evidence on the pass-through of firm-level demand shocks across multiple margins of workplace amenities. Finally, we provide a structural approach that allows us to quantify the degree of substitutability between wages and amenities in worker preferences and the distribution of total compensation markdowns across firms, which we further decompose into a cost-curvature wedge due to the structure of amenity provision costs and a conduct wedge due to employer market power.

Our paper relates to a nascent literature that characterizes how firms determine the optimal amenity-wage bundle in the presence of employer market power (e.g., Dube et al., 2022). We depart from the existing work by allowing multi-dimensional amenities with heterogeneous cost structures and by studying the implications of endogenous amenity provision for the measurement of employer market power.¹ The joint determination of heterogeneous-cost amenities matters because amenities with different cost structures interact through the employment margin: concave-cost amenities amplify the effective labor supply elasticity, while convex-cost amenities dampen it. These interactions have consequences for the measurement of employer market power that would be missed in a single-amenity model. Notably, the composition of the cost structure determines whether the effective labor supply elasticity exceeds or falls short of the structural taste elasticity, which is the underlying source of employer market power. Despite endogenizing the provision of multi-dimensional amenities with heterogeneous costs under employer market power, our framework preserves tractability.

Our work also relates to an empirical literature measuring workplace amenities. Most studies rely on either broad survey categories of benefits or single amenity dimensions such as safety (Lee and Taylor, 2019) or scheduling flexibility (Mas and Pallais, 2017). Mas (2025) highlights the difficulty of capturing the full scope of non-wage job attributes, and Woodbury (1983) and Woodbury and Huang (1991) emphasize that the substitution between wages and benefits depends on the relative prices of many distinct margins, yet most empirical work observes only aggregate fringe-benefit expenditure or a single category. We address this

¹Notably, Mas (2025) considers the case of public-good amenities, which correspond to what we refer to as fixed amenities, and shows in a competitive setting that such amenities create convexity in the firm’s isoprofit curve, making larger firms more likely to provide them. Some recent studies let specific non-pecuniary job attributes also affect output, e.g., work hours in Lachowska et al. (2026), managerial delegation in Di Porto et al. (2026), and schedule flexibility in Adams-Prassl et al. (forthcoming), which we refer to as production-relevant amenities. A related set of studies focuses on how labor market institutions shape the amenity-wage bundle under monopsony, including minimum wages (Dube et al., 2022), union bargaining power (Lagos, forthcoming), and the enforcement of mandated workplace standards (Colonna et al., 2025). We allow multi-dimensional amenities with heterogeneous costs and instead study firm adjustments to demand shocks.

concern using two complementary data sources from Norway. From the administrative employment registers and firm cost accounts, we construct measures of employment conditions (e.g., contract types) and non-wage personnel costs. From the near-universe of posted job ads, we measure workplace amenities as they appear to prospective job seekers, using the text-analysis approach of Audoly et al. (2024). This combination of comprehensive measurement is, to our knowledge, novel and directly addresses the concern that single-dimension studies understate both the firm’s adjustment and the worker’s total valuation response.

Our paper further relates to the literature on the pass-through of firm shocks to workers. The pass-through of firm-level shocks to wages is a central prediction of monopsony models (Kline, 2025), which has been extensively tested (Card et al., 2018; Kline et al., 2019; Lamadon et al., 2022). Some existing evidence suggests that high-paying or more productive firms simultaneously offer better amenities (Caldwell et al., 2025; Roussille and Scuderi, 2025; Sockin, 2022), and that the total compensation gaps between education groups exceed the wage gaps (Maestas et al., 2023). Meanwhile, there is limited direct evidence on how firms adjust amenity provision in response to demand shocks.² Our framework further has implications for estimating firm-level labor supply elasticities. We show that the standard approach to regress log employment on log wages instrumented by a demand shifter identifies an effective supply elasticity that incorporates the endogenous amenity response, as firms simultaneously adjust wages and amenities following the shock, and both channels affect labor supply. This reinterpretation is distinct from standard endogeneity concerns: even with an exogenous instrument, the IV estimand identifies a different object when amenities adjust. We provide conditions under which the total compensation markdowns can be recovered from the estimates of effective labor supply elasticities and amenity cost multipliers, and show that the correction is quantitatively important for measuring employer market power.

Finally, our paper relates to recent work on why employer market power differs across firms. Recent studies focus on the role of preference heterogeneity and worker-firm sorting, through wage-sensitive workers sorting into high-paying firms (Volpe, 2024) or richer workers caring less about wages but sorting into larger firms operating in more concentrated markets (Berger et al., 2026). Alternatively, productive firms may pay above workers’ outside options and thus face less elastic supply (Seegmiller, 2025), while workers with a comparative advantage at their firm have less to gain from leaving (Bils et al., 2025). Our framework instead points to a distinct mechanism in which firm differences reflect the amenity cost structure.

²To best of our knowledge, only two studies have considered such responses to firm-level demand shocks. Kroft et al. (2025) consider firms winning government procurement auctions and report finding no evidence that this affects workers’ safety in the US construction industry (see their Supplemental Appendix H). In contrast, Charles et al. (2022) find that output-price increases in US metal mining raise firm-level injury rates and safety-regulation violations, consistent with firms reallocating effort from safety toward production.

The remainder of the paper proceeds as follows. Section 2 develops the monopsony framework with endogenous, multi-dimensional amenities, derives the equilibrium markdown decomposition, and characterizes firm responses to demand shocks. Section 3 describes the Norwegian administrative data sources, our measurement of workplace amenities and non-wage compensation, the construction of the three firm-level demand shocks, and the event-study research design. Section 4 presents our empirical evidence on multi-margin amenity responses, while Section 5 quantifies the implications for measurement of employer market power. Section 6 concludes with a discussion and avenues for future work.

2 Framework

We provide a monopsony model where firms jointly determine wages and a bundle of amenities. We characterize the equilibrium under endogenous amenity provision and clarify how much of the marginal product accrues to workers depending on the amenity cost structure. Finally, we use the model to derive predictions on firm responses to demand shocks.

2.1 Workers

A population of N workers with utility $U_{ij} = V_j \cdot \xi_{ij}$ chooses among a set of firms $\mathcal{J}$, where V_j is the common utility component and ξ_{ij} are idiosyncratic tastes. The common component is represented by constant elasticity of substitution (CES) preferences over wages w_j and an n -dimensional vector of amenities $\mathbf{a}_j = (a_{j1}, a_{j2}, \dots, a_{jn}) \in \mathbb{R}_+^n$.³ Each worker i thus evaluates the common wage-amenity compensation bundle offered by firm j as follows:

$$V_j \equiv V_j(w_j, \mathbf{a}_j) = \left(w_j^\rho + \sum_{k=1}^n \gamma_k a_{jk}^\rho \right)^{1/\rho}, \quad \rho < 1,$$

where preference weights $\gamma_k > 0$ govern how strongly workers value each amenity k relative to wages, while $\rho < 1$ determines the elasticity of substitution ($1/(1-\rho) > 0$) between wages and amenities.⁴ Wages and amenities can be substitutes ($\rho > 0$) or complements ($\rho < 0$) in the compensation bundle, while in the limiting case $\rho \rightarrow 0$ the common utility index becomes log-additive in wages and amenities (i.e., Cobb–Douglas preferences).

The idiosyncratic tastes are independent draws $\xi_{ij} \sim \text{Fréchet}(\eta, \bar{a}_j)$, where the shape

³For brevity, we abstract from taxation here, but our theoretical results extend easily to the case of a proportional tax on labor earnings. In Appendix A.3, we also consider the case of progressive taxation.

⁴As preference weights $\gamma_k > 0 \forall k$, each component of $\mathbf{a}_j$ enters workers' utility with a non-negative value. Nevertheless, we can accommodate the presence of workplace disamenities (e.g., inconvenient hours or level of injury risk), as a disamenity can be redefined as the deviation from a baseline level of the amenity.

parameter $\eta > 0$ drives the dispersion of the taste shocks, and the firm-specific scale parameter $\bar{a}_j > 0$ captures exogenous differences in firm attractiveness (e.g., location, brand, or industry appeal). Given this structure, we can derive the labor supply facing each firm as:

$$L_j = \frac{N \bar{a}_j}{\mathbb{V}^\eta} V_j^\eta, \quad \mathbb{V}^\eta \equiv \sum_{l \in \mathcal{J}} \bar{a}_l V_l^\eta, \quad (1)$$

where η is the structural labor supply elasticity with respect to the job value V_j and $\mathbb{V}$ is an aggregate attractiveness index over all firms. The formal derivation is in Appendix A.1.1.

2.2 Firms

The firm's output depends only on a homogeneous labor input, $Y(L_j) = \psi_j L_j^\alpha$ with $\alpha \in (0, 1)$, and exogenous productivity $\psi_j > 0$, and the output price is exogenous and normalized to one.⁵ Each firm j is atomistic and takes $\bar{a}_j$ and $\mathbb{V}$ as given, and chooses wages and amenities $(w_j, \mathbf{a}_j)$ to maximize profit subject to its labor-supply curve (1):

$$\max_{w_j, \mathbf{a}_j} \pi_j = \psi_j L_j^\alpha - w_j L_j - \mathbf{C}(\mathbf{a}_j, L_j), \quad \text{s.t.} \quad L_j = \frac{N \bar{a}_j}{\mathbb{V}^\eta} V_j^\eta. \quad (2)$$

The amenity bundle $\mathbf{a}_j$ has a provision cost $\mathbf{C}(\mathbf{a}_j, L_j) \equiv \sum_{k=1}^n C_{jk}(a_{jk}, L_j)$.⁶ We let the cost $C_{jk}(a_{jk}, L_j)$ of amenity k be linear in the level a_{jk} of amenity provision, multiplied by an exogenous firm-specific unit cost factor $c_{jk} > 0$, and a weakly increasing *scaling function* $h_k(L_j) > 0$ governing how the per worker amenity cost depends on firm size, such that:

$$C_{jk}(a_{jk}, L_j) = c_{jk} h_k(L_j) a_{jk}. \quad (3)$$

We parameterize the size dependence through the *size-elasticity of amenity cost*

$$\varepsilon_L^{h_k} \equiv \frac{L_j h'_k(L_j)}{h_k(L_j)},$$

which measures how the per-unit cost of amenity k responds to employment size.

To fix ideas, consider the iso-elastic case $h_k(L_j) = L_j^\delta$ for some $\delta \geq 0$, giving $\varepsilon_L^{h_k} = \delta$. Four canonical regimes span the range from pure “public-good” amenities with fixed costs to amenities with convex costs: (i) *fixed costs* ($\delta = 0$, $h_k = 1$), where per-unit provision

⁵For tractability, we let amenities affect output only through size L_j . In Appendix A.4, we consider the case of *production-relevant* amenities $\{a_{jk}\}_{k \in \mathcal{P}}$ that directly affect output, replacing ψ_j by $\psi_j \prod_{k \in \mathcal{P}} p_k(a_{jk})$.

⁶In Appendix Table A.2, we provide an overview of existing studies of imperfect competition with endogenous amenities, clarifying what assumptions are maintained in these studies regarding the amenity provision cost structure, and how each case compares to the cost structure specified in our framework.

cost is independent of headcount (e.g., an on-site gym or a childcare facility); (ii) *concave-in- L_j costs* ($0 < \delta < 1$, $h_k = L_j^\delta$), capturing provision costs that scale nonlinearly with size due to economies of scale in purchasing or contracting (e.g., bulk discounts on software licenses, negotiated rates for meal plans, group transportation passes, or employee assistance programs); (iii) *per-worker costs* ($\delta = 1$, $h_k = L_j$), where provision costs scale linearly with employment (e.g., pension contributions, or individually paid travel and food compensation for staff); (iv) *convex-in- L_j costs* ($\delta > 1$, $h_k = L_j^\delta$), where provision becomes more costly per worker as the firm grows, due to coordination or congestion (e.g., flexible scheduling that involves pairwise negotiation over shift swaps and meeting alignment, or shared-resource amenities with capacity constraints, such as recreation areas). These regimes span a range of cost structures with different implications for per-worker amenity costs C_{jk}/L_j . Notably, per-worker costs decline (increase) with firm size for concave-cost (convex-cost) amenities.

2.3 Firm Optimum and Equilibrium

We characterize the firm's solution to (2) in Appendix A.1.2, which provides $n + 1$ first-order conditions for wages w_j^* and amenities $\mathbf{a}_j^*$. By combining the solutions for wages and each amenity, we derive the following n MRS conditions:

$$\underbrace{\gamma_k \left(\frac{a_{jk}^*}{w_j^*} \right)^{\rho-1}}_{\text{worker's relative valuation}} = \underbrace{\frac{c_{jk} h_k(L_j)}{L_j}}_{\text{cost per worker}}, \quad \forall k \in \{1, \dots, n\}. \quad (4)$$

The left side is the worker's marginal utility of amenity k relative to wages. The right side is the firm's per-worker cost of a marginal increase in a_{jk} . Rearranging (4), we get

$$\frac{a_{jk}^*}{w_j^*} = \left[\underbrace{\left(\frac{\gamma_k}{c_{jk}} \right)}_{\text{taste-cost ratio}} \cdot \underbrace{\left(\frac{L_j}{h_k(L_j)} \right)}_{\text{cost-scaling factor}} \right]^{\frac{1}{1-\rho}}, \quad \forall k \in \{1, \dots, n\}, \quad (5)$$

where the optimal amenity-wage ratio is pinned down by two forces: (i) The *taste-cost ratio* (γ_k/c_{jk}) captures the worker-firm dimension: higher worker valuation γ_k raises the ratio while higher firm-specific unit cost c_{jk} lowers it, so cross-firm variation in c_{jk} generates variation in amenity provision at any given wage level. (ii) The *cost-scaling factor* ($L_j/h_k(L_j)$) either rises, is invariant, or falls with firm size, as $\varepsilon_L^{h_k} \leq 1$, capturing how amenity provision per unit of wages scales with employment. In all cases, amenity provision is proportional to

wages, scaled by a size-dependent factor that depends on L_j but not on w_j .⁷

The total compensation paid by the firm consists of the wage bill and total amenity costs, $TC_j(w_j, \mathbf{a}_j, L_j) \equiv w_j L_j + \sum_k C_{jk}(a_{jk}, L_j)$. We define the *amenity cost multiplier* as the ratio of total compensation to the wage bill,

$$\mathcal{A}(L_j) \equiv \frac{TC_j(w_j, \mathbf{a}_j, L_j)}{w_j L_j} = 1 + \sum_k \omega_k(L_j), \quad (6)$$

so the firm spends $\mathcal{A}(L_j)$ in total compensation for every unit of spending on the wage bill. Notably, while both equalities follow from the definitions of $\mathcal{A}(L_j)$ and ω_k , what is specific to our model is that the shares depend on L_j alone, where $\omega_k(L_j) \equiv C_{jk}/(w_j L_j)$ is amenity k 's expenditure relative to the wage bill. Although TC_j , $w_j L_j$, and C_{jk} all individually depend on w_j , their ratios do not: substituting the optimal a_{jk}^* from (5) cancels the wage, leaving them as a function of L_j alone, hence the notation $\mathcal{A}(L_j)$ and $\omega_k(L_j)$.⁸ Substituting (5) into ω_k gives

$$\omega_k(L_j) = \left(\frac{\gamma_k}{c_{jk}^\rho} \right)^{\frac{1}{1-\rho}} \left(\frac{L_j}{h_k(L_j)} \right)^{\frac{\rho}{1-\rho}} > 0. \quad (7)$$

Two special cases are worth noting. First, under Cobb–Douglas preferences ($\rho = 0$), we get $\omega_k = \gamma_k$ for all k , implying that amenity expenditure shares simply reflect workers' amenity tastes. Second, if costs scale per-worker ($\varepsilon_L^{h_k} = 1$) for all k , then $\omega_k = (\gamma_k/c_{jk}^\rho)^{1/(1-\rho)}$, so expenditure shares depend on tastes and unit provision costs. In both cases, $\mathcal{A}(L_j)$ is independent of L_j , so the total amenity provision cost is a constant “markup” on the wage.

The MRS conditions (5) pin down each amenity as a wage-proportional function of L_j . Multiplying the MRS condition (4) by a_{jk}/w_j gives $\gamma_k a_{jk}^\rho = \omega_k(L_j) w_j^\rho$, so the wages and amenities contribute w_j^ρ and $\sum_k \omega_k(L_j) w_j^\rho$ to V_j^ρ . The CES aggregator therefore collapses to a wage-proportional job value scaled by the amenity multiplier $\mathcal{A}(L_j) = 1 + \sum_k \omega_k(L_j)$,

$$V_j = \mathcal{A}(L_j)^{1/\rho} w_j. \quad (8)$$

Inserting (8) into the labor-supply curve (1) eliminates the n amenity choices and yields the effective labor supply facing the firm.⁹ Together with the first-order condition for wages, we

⁷For notational ease, we drop the superscript “*” that denotes the firm's optimal solutions in the following discussion of the model's equilibrium, unless we find it necessary to include it to avoid confusion.

⁸The wage-invariance of $\mathcal{A}(L_j)$ and $\omega_k(L_j)$ relies on two features: (i) CES preferences are homothetic, so the worker's MRS depends only on the ratios a_{jk}/w_j ; (ii) amenity provision costs are linear in a_{jk} in (3).

⁹At $\rho = 0$, (8) becomes $V_j \propto w_j \prod_k [L_j/h_k(L_j)]^{\gamma_k/(1+\sum_m \gamma_m)}$, where the constant depends only on the parameters γ_k and c_{jk} . The labor supply equation (10) also holds at $\rho = 0$ as a limit: the factor $\mathcal{A}(L_j)^{\eta/\rho}$ diverges, but so does every term of $\mathbb{V}^\eta$ by the same factor, and their ratio is finite.

characterize the equilibrium of our model by two equations with two unknowns (w_j, L_j) :

$$\text{(Wage Equation)} \quad w_j = \frac{\alpha \psi_j L_j^{\alpha-1}}{1 + \mathcal{A}(L_j)/\eta + \sum_k \omega_k(L_j) \varepsilon_L^{h_k}}, \quad (9)$$

$$\text{(Labor Supply)} \quad L_j = (N \bar{a}_j / \mathbb{V}^\eta) w_j^\eta \mathcal{A}(L_j)^{\eta/\rho}. \quad (10)$$

Solving (9)–(10), we get the equilibrium wage and employment in terms of firm primitives,

$$w_j^* = w^*(\psi_j, \bar{a}_j, \mathbf{c}_j; \mathbb{V}, \mathbf{h}, \boldsymbol{\gamma}, \alpha, \eta, \rho), \quad L_j^* = L^*(\psi_j, \bar{a}_j, \mathbf{c}_j; \mathbb{V}, \mathbf{h}, \boldsymbol{\gamma}, \alpha, \eta, \rho), \quad (11)$$

where $(\mathbf{c}_j, \mathbf{h}, \boldsymbol{\gamma})$ are vectors of firm-specific unit costs, scale-dependent cost functions, and tastes, respectively, associated with the n -dimensional vector of endogenous amenities $\mathbf{a}_j$. The equilibrium amenities $\mathbf{a}_j^*$ follow from evaluating (5) at (w_j^*, L_j^*) , and $\mathcal{A}^* = \mathcal{A}(L_j^*)$. An interior equilibrium exists and is unique when labor supply slopes upward and the wage equation (labor demand) slopes downward (see Appendix A.1.6 for the formal conditions).

The wage equation (9) expresses the wage as a fraction of marginal revenue product (MRP) $\alpha \psi_j L_j^{\alpha-1}$. The denominator comprises the direct wage cost (equal to 1), and two additional amenity-related terms: (i) $\mathcal{A}(L_j)/\eta$ is the infra-marginal compensation cost (raising w_j to attract one more worker also raises every amenity proportionally, since $a_{jk} \propto w_j$), and (ii) $\omega_k(L_j) \varepsilon_L^{h_k}$ is the marginal amenity-provision cost (hiring one more worker raises $h_k(L_j)$ at fixed amenity levels). Both channels widen the wedge between wages and the MRP.¹⁰

Wage-elasticity of the effective labor supply curve. The trade-off that monopsonistic firms face between employment and wages is summarized by the wage-elasticity of the effective labor supply curve. Implicitly differentiating the labor supply equation, $L_j = (N \bar{a}_j / \mathbb{V}^\eta) V_j^\eta$ with $\mathbf{a}_j^*$ substituted out, with respect to the wage yields (Appendix A.1.4):

$$\varepsilon_w^L(L_j) \equiv \left. \frac{d \ln L_j}{d \ln w_j} \right|_{\text{supply}} = \frac{\eta}{1 - \frac{\eta}{1 - \rho} \Phi(L_j)}, \quad \Phi(L_j) \equiv \sum_k \sigma_k(L_j) (1 - \varepsilon_L^{h_k}), \quad (12)$$

where $\sigma_k(L_j) \equiv C_{jk}/TC_j = \omega_k(L_j)/\mathcal{A}(L_j)$ is amenity k 's share of total compensation. The *cost-curvature index* $\Phi(L_j)$ is a σ_k -weighted average of $(1 - \varepsilon_L^{h_k})$: positive when concave-cost amenities dominate, negative when convex-cost amenities dominate, and zero at per-worker costs on aggregate. The effective labor supply elasticity $\varepsilon_w^L(L_j)$ differs from the structural

¹⁰In Appendix Section A.2, we compare the composition of the wage-amenity bundle offered by the monopsonist to the bundle provided under a competitive *value-taking* conduct. Proposition A.1 shows that the monopsonist will strictly under-provide employment, wages, and job value as compared to the value-taking benchmark, while amenities will be under-provided unless the provision costs are strongly convex.

taste dispersion parameter η as the amenity bundle also adjusts when w_j changes and workers respond to the complete change in the wage-amenity compensation bundle. Specifically, if $\Phi(L_j) > 0$, the amenity adjustment amplifies the supply response ($\varepsilon_w^L(L_j) > \eta$), while if $\Phi(L_j) < 0$, the amenity adjustment dampens the supply response ($\varepsilon_w^L(L_j) < \eta$).¹¹

In recent work, Kline (2025) points out that common iso-elastic supply specifications used in the monopsony literature impose a zero super-elasticity. Notably, as $\Phi(L_j)$ itself depends on firm size, $\varepsilon_w^L(L_j)$ is not constant along the supply curve. The super-elasticity is generally non-zero in our framework, and it depends crucially on the cost structure (A.7).

Size-elasticity of the amenity cost multiplier. How amenity provision costs vary with firm size is captured by the size-elasticity of the amenity cost multiplier. Differentiating the multiplier $\mathcal{A}(L_j) = 1 + \sum_k \omega_k(L_j)$ along the labor supply curve (Appendix A.1.4), gives its size-elasticity as the σ_k -weighted counterpart of the cost-curvature index,

$$\varepsilon_L^{\mathcal{A}}(L_j) \equiv \frac{d \ln \mathcal{A}(L_j)}{d \ln L_j} = \sum_k \sigma_k(L_j) \frac{\rho}{1 - \rho} (1 - \varepsilon_L^{h_k}) = \frac{\rho}{1 - \rho} \Phi(L_j). \quad (13)$$

It captures how amenity spending per wage unit scales with employment: the multiplier rises with firm size when ρ and $\Phi(L_j)$ share sign, falls when they oppose, and is invariant ($\varepsilon_L^{\mathcal{A}} = 0$) under per-worker costs on aggregate ($\Phi(L_j) = 0$) or under Cobb–Douglas preferences ($\rho = 0$). Notably, if $\rho = 0$, then the amenity cost multiplier is constant ($\varepsilon_L^{\mathcal{A}} = 0$) for *any* cost structure, since fixed expenditure shares keep spending per wage unit constant even as amenity *levels* move. Although $\rho = 0$ implies $\varepsilon_L^{\mathcal{A}} = 0$, the cost-curvature index $\Phi(L_j)$ still shifts the supply elasticity (12), as workers respond to the amenity levels that are provided with the wage.

2.4 Markdowns and Amenity Cost Structure

How much of the marginal product of labor accrues to workers in equilibrium? We derive the total compensation markdown from the firm’s optimum, where its marginal cost equals its marginal product. Differentiating the compensation bill $L_j \mathcal{A}(L_j) w_j$ along the supply curve,

$$\text{MC}(L_j) \equiv \frac{d[L_j \mathcal{A}(L_j) w_j]}{dL_j} = \mathcal{A}(L_j) w_j + L_j \frac{d[\mathcal{A}(L_j) w_j]}{dL_j},$$

¹¹Note that if $\eta \Phi(L_j) > 1 - \rho$, the denominator of (12) would turn negative and the labor supply curve would bend backwards. We throughout maintain the assumption that $\eta \Phi(L_j) < 1 - \rho$ (Condition (i) of Appendix A.1.6), which effectively places an upper bound on the degree of concavity of the cost-curvature index. In Section 5, we show that the effective labor supply elasticity estimates are indeed strictly positive.

where substituting the *total compensation supply elasticity* $\varepsilon_{\mathcal{A}w}^L$, the response of L_j to total compensation $\mathcal{A}(L_j)w_j$ along the supply curve, gives:

$$\begin{aligned} \text{MC}(L_j) &= \mathcal{A}(L_j) w_j \left(1 + \underbrace{\frac{d \ln[\mathcal{A}(L_j)w_j]}{d \ln L_j}}_{\equiv (\varepsilon_{\mathcal{A}w}^L)^{-1}} \right) = \mathcal{A}(L_j) w_j \left(1 + \underbrace{\frac{d \ln w_j}{d \ln L_j}}_{=(\varepsilon_w^L)^{-1}} + \underbrace{\frac{d \ln \mathcal{A}(L_j)}{d \ln L_j}}_{=\varepsilon_L^A} \right) \\ &= \mathcal{A}(L_j) w_j \left(1 + \frac{1}{\varepsilon_{\mathcal{A}w}^L} \right) = \mathcal{A}(L_j) w_j \left(1 + \frac{1}{\varepsilon_w^L} + \varepsilon_L^A \right). \end{aligned}$$

Equating $\text{MC}(L_j) = \text{MRP}(L_j)$, rearranging and defining $\mathcal{M}(L_j) \equiv w_j/\text{MRP}(L_j)$, yields:

$$\mathcal{A}(L_j)\mathcal{M}(L_j) = \frac{\varepsilon_{\mathcal{A}w}^L}{1 + \varepsilon_{\mathcal{A}w}^L} = \frac{\varepsilon_w^L}{(\varepsilon_w^L + 1) + \varepsilon_w^L \varepsilon_L^A} = \frac{\eta}{\eta + 1 - \eta \Phi(L_j)}, \quad (14)$$

where we refer to $\mathcal{A}(L_j)\mathcal{M}(L_j)$ as the *total compensation markdown*. Equation (14) provides three equivalent expressions for this markdown. The first expression is the canonical inverse-elasticity monopsony markdown formula applied to $\varepsilon_{\mathcal{A}w}^L$. The second decomposes it through ε_w^L and ε_L^A ; the wedge $\varepsilon_w^L \varepsilon_L^A$ is the amenity-feedback correction relative to the canonical markdown formula $\varepsilon_w^L/(1 + \varepsilon_w^L)$. The third expresses it in terms of the Fréchet-shape parameter η , which represents the structural elasticity (reflecting taste dispersion) in the canonical monopsony model, and $\Phi(L_j)$, which captures the amenity cost structure.¹²

Dividing (14) by $\mathcal{A}(L_j)$, we can further derive the canonical *wage markdown* $\mathcal{M}(L_j) \equiv w_j/\text{MRP}(L_j)$ as the total compensation markdown deflated by the amenity multiplier:

$$\mathcal{M}(L_j) = \underbrace{\frac{1}{\mathcal{A}(L_j)}}_{\text{amenity deflator}} \cdot \frac{\eta}{\eta + 1 - \underbrace{\eta \Phi(L_j)}_{\text{cost curvature}}}. \quad (15)$$

Three observations follow. First, the equilibrium total compensation markdown $\mathcal{A}^*\mathcal{M}^*$ deviates from the canonical $\eta/(\eta + 1)$ benchmark in the direction of $\Phi(L_j)$: the total compensation markdown is above the canonical markdown when $\Phi(L_j) > 0$ (concave costs), below when $\Phi(L_j) < 0$ (convex costs), and at the benchmark when $\Phi(L_j) = 0$ (per-worker costs). Second, $\mathcal{M}(L_j)$ sits below $\mathcal{A}(L_j)\mathcal{M}(L_j)$ by the factor $\mathcal{A}(L_j) > 1$. This wedge reflects relabeling between wages and amenities, not between workers and the firm. Third, the first expression in (14) is the canonical inverse-elasticity formula applied to the *total compensation* elasticity $\varepsilon_{\mathcal{A}w}^L$. It reduces to the wage-based formula when the two elasticities coincide, that is, when the amenity cost multiplier is invariant ($\varepsilon_L^A = 0$, since $1/\varepsilon_{\mathcal{A}w}^L = 1/\varepsilon_w^L + \varepsilon_L^A$).

¹²To get the third expression, we plug in the closed forms of ε_w^L and ε_L^A from (12) and (13), or equivalently $\varepsilon_{\mathcal{A}w}^L = \eta/[1 - \eta \Phi(L_j)]$ obtained via the identity $1/\varepsilon_{\mathcal{A}w}^L = 1/\varepsilon_w^L + \varepsilon_L^A$. The same substitution recovers (A.13).

By (13) this holds under Cobb–Douglas preferences ($\rho = 0$) or aggregate per-worker costs ($\Phi(L_j) = 0$). In either case, $\mathcal{A}^* = \bar{\mathcal{A}}$, and thus we have $\mathcal{A}^*\mathcal{M}^* = \bar{\mathcal{A}}\mathcal{M}^* = \varepsilon_w^L/(\varepsilon_w^L + 1)$, so wage-based elasticity measures can recover the total compensation markdown even though wages are only part of the compensation. We summarize these observations in Proposition 1.


Proposition 1 (Markdowns and Amenity Cost Structure). *With any valued endogenous amenity: (i) The total compensation markdown deviates from the canonical benchmark in the direction of the cost-curvature index: $\mathcal{A}(L_j)\mathcal{M}(L_j) \gtrless \eta/(\eta + 1)$ iff $\Phi(L_j) \gtrless 0$, with equality at the per-worker benchmark $\Phi(L_j) = 0$. (ii) The wage markdown $\mathcal{M}(L_j)$ is strictly below the total compensation markdown $\mathcal{A}(L_j)\mathcal{M}(L_j)$ by the factor $\mathcal{A}(L_j) > 1$. (iii) The canonical formula recovers the total compensation markdown exactly, $\mathcal{A}^*\mathcal{M}^* = \varepsilon_w^L/(\varepsilon_w^L + 1)$, in two special cases: (a) Cobb–Douglas preferences ($\rho = 0$), and (b) per-worker amenity costs on aggregate ($\Phi(L_j) = 0$). In both, the amenity cost multiplier is constant, $\mathcal{A}^* = \bar{\mathcal{A}}$, so the wage markdown is $\mathcal{M}^* = \mathcal{A}^*\mathcal{M}^*/\bar{\mathcal{A}}$. In case (b), in addition, $\varepsilon_w^L = \eta$.*

Figure 1 maps the markdown decomposition in the (L_j, w_j) space. In canonical monopoly without endogenous amenities (panel a), $\text{MC}(L_j)$ crosses $\text{MRP}(L_j)$ at an employment level below the competitive benchmark, and the wage markdown wedge $1 - \mathcal{M}^*$ separates the wage from the marginal product. With endogenous amenities (panels b–d), total compensation $\mathcal{A}(L_j)w_j$ replaces w_j as the firm’s effective price of labor and the total compensation markdown $\mathcal{A}^*\mathcal{M}^*$ is the welfare-relevant object. At the optimum, the marginal revenue product MRP^* splits into three parts: wage w_j^* , amenity cost $(\mathcal{A}^* - 1)w_j^*$, and firm surplus $(1 - \mathcal{A}^*\mathcal{M}^*)\text{MRP}^*$. The wedge between the wage and the marginal product, $1 - \mathcal{M}^*$, widens as amenities absorb part of what would otherwise be wages. The three cost regimes share a common mechanism: hiring one more worker triggers a re-optimization of amenity provision, which feeds back into the supply curve. Whether the feedback amplifies or dampens supply, and hence whether $\mathcal{A}^*\mathcal{M}^*$ exceeds or falls short of the canonical $\eta/(\eta + 1)$, depends on how average amenity cost per worker scales with firm-size.¹³


We now review each cost regime in turn: (i) With per-worker costs ($\Phi(L_j) = 0$), amenity spending is scale-invariant, so there is a *neutral feedback* on labor supply: $\varepsilon_w^L = \eta$. We have $\text{MC} = \mathcal{A}(L_j)w_j(1 + 1/\eta)$, which mirrors the canonical $\text{MC} = w_j(1 + 1/\eta)$, so $\mathcal{A}^*\mathcal{M}^* = \eta/(\eta + 1)$. Panel (b) therefore reproduces panel (a) with the price of labor relabeled from w to $\mathcal{A}(L_j)w$, and the same equilibrium and competitive employment. (ii) Under concave amenity costs ($\Phi(L_j) > 0$), amenity spending is diluted by firm size, as fixed expenses spread across more workers, so we get *positive feedback* on labor supply. The firm provides *more*

¹³The illustration sets $\rho = 0$, so $\mathcal{A}(L_j)$ is constant within each panel and the panels differ only through $\Phi(L_j)$, which moves the supply elasticity ε_w^L and the firm-surplus wedge together.


Figure 1: Markdowns across Alternative Amenity-Cost Regimes.




(a) Canonical Monopsony.



(b) Per-Worker Amenity Costs ($\Phi(L_j) = 0$).



(c) Concave Amenity Costs ($\Phi(L_j) > 0$).



(d) Convex Amenity Costs ($\Phi(L_j) < 0$).

Notes: Each panel plots the firm's choice in (L_j, w_j) space, with the equilibrium (L_j^*, w_j^*) marked by black dashed lines. Panel (a) is the canonical benchmark without endogenous amenities: the wedge $1 - \mathcal{M}^*$ separates w_j^* from MRP^* , and the markdown equals the canonical formula $\eta/(\eta+1)$. Panels (b)–(d) introduce endogenous amenities, so MRP^* splits three ways—wage w_j^* , amenity cost $(\mathcal{A}^* - 1)w_j^*$, and firm surplus $(1 - \mathcal{A}^*\mathcal{M}^*)MRP^*$. At per-worker costs (b), $\mathcal{A}(L_j)$ is constant in L_j , and the total compensation markdown $\mathcal{A}^*\mathcal{M}^*$ equals the canonical $\eta/(\eta+1)$. Concave costs (c) raise $\mathcal{A}^*\mathcal{M}^*$ above the benchmark, narrowing firm surplus; convex costs (d) lower it, widening firm surplus. The illustration sets $\rho = 0$, so $\mathcal{A}(L_j)$ is constant in L_j ($\varepsilon_L^A = 0$). Different cost technologies generate different $\Phi(L_j)$ values, which shift the equilibrium supply elasticity ε_w^L and the firm-surplus wedge $1 - \mathcal{A}^*\mathcal{M}^*$.

amenity per worker as it grows. This attracts workers at any wage, which flattens $\mathcal{A}(L_j)w$ and lifts ε_w^L above η . The total compensation markdown $\mathcal{A}^*\mathcal{M}^*$ rises above $\eta/(\eta+1)$ and the wedge at the firm's optimum narrows. Equilibrium employment is higher than at per-worker costs (panel c). (iii) Under convex amenity costs ($\Phi(L_j) < 0$), each additional worker is

more expensive to equip, so amenities fall relative to the wage as the firm grows, and we get *negative feedback* on labor supply: attracting the next worker requires a larger wage hike, which steepens $\mathcal{A}(L_j)w$ and pulls ε_w^L below η . The total compensation markdown $\mathcal{A}^*\mathcal{M}^*$ falls below $\eta/(\eta + 1)$ and the wedge at the firm's optimum widens, so the intuition that “amenities reduce markdowns” reverses. Equilibrium employment is lower (panel d).

Size-elasticity of the wage markdowns. How wage markdowns adjust with firm size is captured by the elasticity $\varepsilon_L^{\mathcal{M}}(L_j) \equiv \frac{d \ln \mathcal{M}(L_j)}{d \ln L_j}$. We derive the functional form of this elasticity in Appendix A.1.5. Notably, under iso-elastic costs and substitutes ($\rho > 0$), $\varepsilon_L^{\mathcal{M}} < 0$ if the firm faces concave costs, so markdowns deepen as firms grow. And, $\varepsilon_L^{\mathcal{M}} > 0$ if the firm faces convex costs, so markdowns are more compressed in larger firms. The wage markdown's size-elasticity will partly determine how firm size mediates the wage and employment responses to productivity shocks in our comparative statics in Section 2.5.

Conduct versus cost structure. Having defined the total compensation markdown as the gap between total compensation and marginal product in (14), we ask how much of this gap reflects employer market power, and how much the structure of amenity provision costs. We can express the total compensation markdown as the product of two factors:

$$\mathcal{A}(L_j) \mathcal{M}(L_j) = \underbrace{\frac{1}{1 - \Phi(L_j)}}_{\text{cost curvature}} \times \underbrace{\frac{1 - \Phi(L_j)}{1 - \Phi(L_j) + 1/\eta}}_{\text{conduct}}. \quad (16)$$

This decomposition serves as an *accounting identity* featuring two terms that can be interpreted as capturing the cost curvature wedge and the conduct wedge, respectively. Intuitively, while the second factor tends to one as $\eta \rightarrow \infty$, a firm without employer market power would nevertheless pay $1/(1 - \Phi(L_j))$ of its marginal product. Attributing the total compensation markdown as a whole to employer market power would conflate cost curvature with conduct. In Section 5, we thus decompose the cost curvature and conduct components.¹⁴

2.5 Productivity Shocks

We now characterize how the firm's wages, employment, and amenity provision respond to a positive firm-specific productivity shock $d\psi_j > 0$. In our framework, the comparative statics for a negative productivity shock, $d\psi_j < 0$, are symmetric, and the predictions further

¹⁴We interpret this decomposition as an accounting exercise as it is valid only at the monopsonist's employment level L_j^* . A value-taking firm, whose marginal cost (A.12) lacks the conduct term $\mathcal{A}(L_j)w(L_j)/\eta$ of the monopsonist's marginal cost (A.13), would expand to $L_j^c > L_j^*$ (Proposition A.1), where its compensation is $1/(1 - \Phi(L_j^c))$ of its marginal product and the cost-curvature wedge differs from the one at L_j^* .

carry over to the same-signed firm-level product-demand shocks.¹⁵ The productivity shock enters the equilibrium conditions of our model only through the wage equation (9). Total log-differentiation of the wage equation (9) and labor supply (10) yields

$$\begin{aligned}\text{Wage Equation: } d \ln w_j &= d \ln \psi_j - [(1 - \alpha) - \varepsilon_L^M] d \ln L_j, \\ \text{Labor Supply: } d \ln L_j &= \varepsilon_w^L d \ln w_j,\end{aligned}$$

where the bracketed term is the wage-equation slope, ε_L^M is the wage markdown's size-elasticity from (A.8), and ε_w^L is the labor supply elasticity from (12). Substituting gives

$$\frac{d \ln w_j}{d \ln \psi_j} = \frac{1}{1 + \varepsilon_w^L(1 - \alpha - \varepsilon_L^M)} > 0, \quad \frac{d \ln L_j}{d \ln \psi_j} = \varepsilon_w^L \cdot \frac{d \ln w_j}{d \ln \psi_j} > 0,$$

where the denominator is positive under Conditions (i) and (ii) of Appendix A.1.6, so both wages and employment rise. The ratio of these two responses to a productivity shifter is the natural empirical analog of the *local* labor supply elasticity, $d \ln L_j / d \ln w_j = \varepsilon_w^L(L_j)$, and underlies the instrumental variables (IV) strategy in Lamadon et al. (2022).¹⁶

How amenities respond is more nuanced and depends on what is measured: the amenity level a_{jk} and the per-worker amenity spending C_{jk}/L_j have distinct comparative statics, and the two need not move together. Since a_{jk}^* is wage-proportional with a size-dependent factor, log-differentiating (5) with respect to ψ_j and using $d \ln(L_j/h_k(L_j))/d \ln L_j = 1 - \varepsilon_L^{h_k}$ decomposes the level response into a direct wage channel and an indirect firm-size channel:

$$\frac{d \ln a_{jk}}{d \ln \psi_j} = \left[\underbrace{1}_{\text{wage channel}} + \underbrace{\frac{1}{1 - \rho}(1 - \varepsilon_L^{h_k})\varepsilon_w^L}_{\text{firm-size channel}} \right] \frac{d \ln w_j}{d \ln \psi_j}. \quad (17)$$

The wage channel captures that wages and amenities must always move together according to the MRS condition (5) when the firm-size channel is absent. The sign of the firm-size channel depends on $1 - \varepsilon_L^{h_k}$: positive at concave costs ($\varepsilon_L^{h_k} < 1$), zero at per-worker costs ($\varepsilon_L^{h_k} = 1$), negative at convex costs ($\varepsilon_L^{h_k} > 1$), while its strength depends on $\varepsilon_w^L(L_j)$ and ρ .¹⁷

Per-worker spending on amenity k is $C_{jk}/L_j = \omega_k(L_j) w_j$, the wage scaled by ω_k from (7),

¹⁵We normalized product prices $p_j = 1$ in Section 2.2. Under firm-level price heterogeneity p_j , the comparative statics of a product price change dp_j are identical to those of a productivity change $d\psi_j$.

¹⁶As $\varepsilon_w^L(L_j)$ varies with firm size when the super-elasticity (A.7) is non-zero, an IV regression that pools across firms produces a weighted average of local elasticities rather than a single structural target parameter. The IV estimand collapses to a constant target parameter in two special cases: (a) Cobb–Douglas preferences ($\rho = 0$) combined with iso-elastic costs $\varepsilon_L^{h_k} = \delta \geq 0 \forall k$, reducing the cost-curvature index to a constant ($\Phi(L_j) = \bar{\Phi}$); (b) per-worker amenity costs scale on aggregate ($\Phi(L_j) = 0$), where $\varepsilon_w^L = \eta$ exactly.

¹⁷The bracket in (17) equals ε_w^L times the supply-curve gradient (A.14) from Proposition A.1: the shock moves the firm up its supply curve, so the amenity level rises precisely when the gradient (A.14) is positive.

which is amenity k 's expenditure relative to the wage bill. Log-differentiating, $d \ln(C_{jk}/L_j) = d \ln \omega_k(L_j) + d \ln w_j$, and substituting $d \ln \omega_k(L_j)/d \ln L_j = \rho(1 - \varepsilon_L^{h_k})/(1 - \rho)$ from (A.5), together with $d \ln L_j = \varepsilon_w^L d \ln w_j$ gives

$$\frac{d \ln(C_{jk}/L_j)}{d \ln \psi_j} = \left[\underbrace{1}_{\text{wage channel}} + \underbrace{\frac{\rho}{1 - \rho}(1 - \varepsilon_L^{h_k})\varepsilon_w^L}_{\text{share channel}} \right] \frac{d \ln w_j}{d \ln \psi_j}. \quad (18)$$

The structure of spending responses mirrors (17) except the second channel carries an extra factor ρ . Under Cobb–Douglas preferences ($\rho = 0$) or per-worker amenity cost structure ($\varepsilon_L^{h_k} = 1$), the share channel vanishes and the adjustment in per-worker amenity spending tracks wages exactly. Otherwise, the share channel's sign depends on $\rho(1 - \varepsilon_L^{h_k})$: it adds to the wage channel under substitutes with concave costs ($\rho > 0, \varepsilon_L^{h_k} < 1$) or complements with convex costs ($\rho < 0, \varepsilon_L^{h_k} > 1$), and subtracts under the opposite pairings. Thus, per-worker spending falls only when the share channel is negative and outweighs the wage channel.

Taken together, we note that the sign of the amenity level response depends on whether the cost structure is sufficiently convex, while the per-worker spending response depends on both the degree of cost convexity *and* the wage-amenity substitutability or complementarity in preferences (the sign of ρ). The two responses can therefore move in opposite directions: substitutes with strongly convex costs can deliver a falling amenity level alongside rising per-worker spending, while complements with strongly concave costs can deliver the reverse.¹⁸

The bundle aggregates these individual responses through the amenity multiplier $\mathcal{A}(L_j)$. Using the expression for the multiplier size-elasticity $\varepsilon_L^{\mathcal{A}}$ from (13) and $d \ln L_j = \varepsilon_w^L d \ln w_j$,

$$\frac{d \ln \mathcal{A}(L_j)}{d \ln \psi_j} = \varepsilon_L^{\mathcal{A}} \cdot \varepsilon_w^L \cdot \frac{d \ln w_j}{d \ln \psi_j}. \quad (19)$$

Thus, the sign of the amenity cost multiplier response is governed by the sign of the multiplier size-elasticity, as the last two terms are positive. Analogous to the share channel for per-worker amenity spending, the sign of $\varepsilon_L^{\mathcal{A}}$ depends jointly on the preference structure ρ and the cost curvature index $\Phi(L_j)$. Crucially, it rises only when $\text{sgn}(\rho) = \text{sgn}(\Phi(L_j))$. Collecting the sign conditions for the three objects yields the following comparative statics.

Proposition 2 (Productivity and Endogenous Amenity Provision). *Following a positive productivity shock $d\psi_j > 0$: (i) The amenity level a_{jk}^* rises if and only if $\varepsilon_L^{h_k} - 1 < \frac{1-\rho}{\varepsilon_w^L}$. (ii) Per-worker amenity spending C_{jk}^*/L_j^* rises if and only if $\rho(\varepsilon_L^{h_k} - 1) < \frac{1-\rho}{\varepsilon_w^L}$, which fails*

¹⁸Notably, these comparative statics maintain that workplace amenities are production-neutral. In Appendix A.4, we consider the case of production-relevant amenities and show that amenities paid largely in forgone output can decline after a positive shock even if their monetary cost is purely per-worker.

if ρ and $\varepsilon_L^{h_k} - 1$ share sign and their product exceeds the threshold: substitutes with strongly convex costs, or complements with strongly concave costs. **(iii)** The amenity cost multiplier $\mathcal{A}^*$ rises if and only if ρ and $\Phi(L_j)$ share sign, and is invariant at $\rho = 0$ or $\Phi(L_j) = 0$.

Empirically, the level response of an amenity relative to the wage response isolates the sign of its cost curvature, whereas both per-worker amenity spending and the amenity multiplier combine preference patterns with cost curvature.¹⁹ Their direction, therefore, identifies the sign of the interaction between these elements rather than either element in isolation. Interpreting spending responses requires having knowledge about the pattern of preferences.

2.6 Identification of Structural Parameters

The comparative statics above guide the identification of η and ρ in our model. A productivity shock $d\psi_j$ traces out, simultaneously, the effective labor supply elasticity $\varepsilon_w^L = d\ln L_j/d\ln w_j$ and the amenity cost multiplier elasticity $\varepsilon_L^A = d\ln \mathcal{A}(L_j)/d\ln L_j$. Combining (12) and (13), we can derive the structural relationship between these objects:

$$\frac{1}{\varepsilon_w^L} = \frac{1}{\eta} - \frac{1}{\rho} \varepsilon_L^A. \quad (20)$$

As η and ρ are common to all firms while $\Phi(L_j)$ varies with size, the pair $(\varepsilon_w^L, \varepsilon_L^A)$ traced across the distribution of firms identifies both of these structural parameters: across firms, $1/\varepsilon_w^L$ moves with ε_L^A at slope $-1/\rho$ and intercept $1/\eta$. With (η, ρ) in hand, the cost-curvature index $\Phi(L_j)$ can then be recovered from (12) or (13). Appendix B formalizes this identification argument when $\Phi(L_j)$ varies with size and $\rho \neq 0$, and discusses the range of parameter values that can be set-identified under alternative scenarios when $\Phi(L_j)$ is degenerate or $\rho = 0$. In the Cobb–Douglas case, the level of $\Phi(L_j)$ is not point-identified, but Appendix B.5 shows that a sign restriction on cost elasticities then bounds the decomposition of the markdown.

¹⁹We illustrate the implications of Proposition 2 numerically in Appendix A.6. Specifically, we trace the behavior of six equilibrium objects—wages, amenity multiplier, labor (wage bill) share, total compensation share, wage markdown and total compensation markdown—for the stylized case of a single amenity under alternative calibrations of amenity cost structures and wage-amenity preferences in our model. The amenity cost multiplier behaves as the proposition predicts: at Cobb–Douglas ($\rho = 0$, Appendix Figure A.2), the multiplier is constant, with $\mathcal{A}(L_j)$ invariant in L_j regardless of $\Phi(L_j)$, while with substitutes ($\rho > 0$, Appendix Figure A.3) and complements ($\rho < 0$, Appendix Figure A.4), the cost structure determines its shape. Concave costs lift the amenity cost multiplier under substitutes and depress it under complements, while convex costs do the reverse. Unless the cost-curvature index is constant ($\Phi(L_j) = \bar{\Phi}$), our model predicts that total compensation markdowns and total compensation (labor) shares vary across the productivity distribution, with profiles that depend both on the structure of amenity costs and preferences. Wage markdowns vary unless, in addition, preferences are Cobb–Douglas ($\rho = 0$) or costs are per-worker ($\Phi(L_j) = 0$).

3 Data and Empirical Strategy

We describe our data sources, including measures of workplace amenities and three alternative firm-level demand shocks, and present our empirical strategy. Our analysis combines several Norwegian administrative data sources that capture standard firm- and worker-level outcomes, along with non-pecuniary job attributes recorded in employment registers, detailed firm cost accounts that capture non-wage compensation, and job ads that provide information on amenities offered to prospective hires. We focus on private-sector firms throughout. Our data sources span the period 2002–2022. We focus on the period 2015–2022 in some of our analyses, where we have more detailed information on contract types.

3.1 Measurement of Amenities and Non-Wage Compensation

We use three distinct data sources for the measurement of workplace amenities and non-wage compensation: (i) monthly matched employer–employee data, which provide worker- and employer-level information on contract types and in-kind compensation, (ii) annual firm-level cost accounts, which provide information on non-wage compensation, and (iii) a near-universe of publicly posted job vacancies, including the full-text portion of postings, with information on advertised workplace amenities and other job attributes at the ad level.

We have access to population-wide, matched employer–employee records linking all workers to their employers, from which we construct annual firm-level employment measures and annual worker- and firm-level hourly wages for the period 2002–2022. We also use more detailed monthly records (A-meldingen), which provide worker-employer-level information on contract type: from January 2015, we can define indicators for full-time or part-time contracts, fixed-hours or variable-hours contracts, and jobs involving shift work, and from January 2021, we observe whether a job is permanent or temporary. For our analyses, we construct annual firm-level measures of shares of full-time, fixed-hours (non-variable), permanent, and regular (non-shift-work) contracts.²⁰ Further, we also have information on the monetary values of compensation for in-kind goods received in addition to wages, such as access to a free car, covered electronic communication (e.g., cellphone), covered lodging or food, employee loans, covered personal membership at a fitness center, or travel allowances, which are all reported by the employer as part of the taxable in-kind compensation to the authorities. We use this information to construct annual firm-level measures of in-kind compensation per worker and of the share of workers receiving in-kind compensation.

Using the information drawn from the matched employer–employee records—contract types and in-kind compensation—we can construct annual firm-level measures, respectively,

²⁰To abstract from seasonal patterns, our reference month with A-meldingen is November.

across all workers employed in the firm, for stayers who remained in the same job across multiple years, and for incumbents who worked in a firm for some years and later changed to a different employer in our data. To avoid concerns about compositional changes in the workforce following demand shocks, we primarily focus on the stayers sample—workers who remained employed at the same firm throughout the event window.²¹

We also have access to annual firm-level cost accounts, which allow us to decompose total labor costs into wage bill and non-wage personnel costs. These data are available for the period 2002–2022. Non-wage personnel costs consist of two components: employer-sponsored pension contributions above the regulated minimum and other non-wage personnel costs (e.g., employee perks, gifts, gym access, company cabin, travel and food allowances).²² Notably, taxable non-wage personnel costs that are divisible across workers, such as employee perks, are also reported in the employment records and contained in our measures of in-kind compensation at the worker-firm level. However, non-wage personnel costs at the firm level include additional components beyond divisible in-kind compensation, such as employer-sponsored pension contributions, as well as non-divisible travel and food-related expenses (e.g., canteen costs). Having observable firm-level non-wage labor costs allows us to quantify measures of total compensation costs beyond hourly wages. Using firm-level employment counts, we also construct per-worker measures for the non-wage labor cost variables.

As an additional source of information on workplace amenities, we also use the near-universe of Norwegian job postings. Specifically, we use information on workplace attributes that workers may value and that are publicly posted in job ads by employers but not necessarily reflected in firms’ non-wage personnel costs. Such attributes could be salient aspects of employment conditions that are valuable to workers (e.g., the possibility to work from home or flexible hours) but not captured in our measures of contract types drawn from the matched employer–employee records, or reflect more subjective aspects of the workplace (e.g., a good social environment), as communicated by the employers in job postings. Recently, Audoly et al. (2024) developed a text-analysis methodology to extract such information about workplace amenities from the full text of job postings. We use their approach to assign each job ad with mentions of specific workplace amenities and further use information on the employer posting each ad to construct annual firm-level measures of the share of postings listing each workplace amenity for our analysis. Our data sources include a register of job ads

²¹We applied initial restrictions to the dataset, keeping workers with a non-missing firm identifier and at least 52×4 annual hours of work, and retained the highest-wage contract per worker-firm-year.

²²The annual firm-level cost accounts also provide information on payroll taxes, covering public pensions, sickness leave payments, unemployment insurance, etc. However, as payroll taxes are statutorily determined and levied proportionally to wage costs, there is no scope for firm-level adjustments in these (relative to wage costs) in response to shocks. We thus do not include payroll taxes in our measures of non-wage compensation.

maintained by the Norwegian Labour and Welfare Administration (NAV) and a collection of job ads posted on the major private online job platform FINN.²³ Each job posting is categorized according to mentions of amenity attributes that we classify into different categories: the presence of on-the-job training, the presence of financial benefits (such as pensions, insurance, or mortgage support), workplace quality (such as interesting tasks, remote work, or social environment), additional perks (such as gym, meals, parking, or company car), inconvenient working hours (such as shift work), flexible working hours, and good career opportunities. These measures capture amenities as they appear to prospective hires in job ads, which can be thought of as providing valuable signals about what workers may expect the firm to offer upon hiring.²⁴

Appendix Table C.1 provides an overview of labor cost components and amenity measures we use, while Appendix Figure C.4 shows the prevalence of non-wage compensation and non-pecuniary job attributes across the firm value-added distribution. High value-added firms have higher hourly wage costs as well as higher hourly non-wage compensation, and most non-pecuniary job attributes vary systematically across the firm distribution. In our data, the ratio of non-wage compensation to the wage bill varies around 10% across firms, with high value-added firms having a higher share of compensation on pension premiums.

To summarize, while the firm cost accounts data provide precise information on the monetary cost of amenity provision, allowing us to directly measure how firms allocate resources across margins of non-wage compensation, information on non-wage job attributes from matched employer–employee data and job postings allows us to capture measures of amenities that are not necessarily visible in firms’ non-wage personnel costs. Notably, the measures of non-pecuniary job attributes and non-wage compensation derived from high-quality administrative registers have limited scope for measurement error. Moreover, having information on both wages and per-worker measures of non-wage personnel costs allows us to link our empirical analysis to the theoretical insights developed above. Since different non-wage components likely scale their provision costs with firm size differently—pension

²³As detailed in Audoly et al. (2024), the employment agency collects information about vacancies from several sources, including online job boards and newspapers, as well as direct job opening reports from employers. Under the Labor Market Act §7, Norwegian employers must report publicly posted job openings to the agency, which maintains a comprehensive database of such ads to provide job-seekers with current information on suitable job opportunities. We accessed similar information from the FINN job platform for our analysis. Both sources provide employer identifiers, which allow us to construct firm-level measures of offered workplace amenities and to link job ad data to our remaining data sources. As in Audoly et al. (2024), we exclude job ads posted by recruitment or temporary employment agencies, as we cannot link such ads to the actual firm or establishment where the job is posted.

²⁴Notably, mentions of amenities detected in job ads are highly predictive of revealed-preference measures of job quality, such as the employer value of Sorkin (2018) and the poaching index of Bagger and Lentz (2019), and validated against realized attributes and choices from a survey experiment (Audoly et al., 2024).

premiums may scale per worker, while items such as gym facilities or company cabins involve fixed costs—disaggregating non-wage components allows us to empirically test the model’s predictions across amenities with potentially different cost structures. As discussed in Section 2.5, our model makes distinct predictions about how demand shocks affect the level of amenity provision and per-worker amenity spending. Having measures of both allows drawing sharper inference. The job ads data complement the cost-based and tangible measures of contract types by capturing whether firms advertise more subjective, harder-to-quantify attributes, allowing us to track adjustments across a comprehensive amenity bundle. Together, the three sources provide the most comprehensive multi-margin coverage of amenities we are aware of. This addresses an important concern in Mas (2025) that studies focusing on a single amenity dimension understate the firm’s potential total adjustment.

3.2 Measurement of Demand Shocks

We use three alternative measures of firm-level demand shocks. First, we use annual firm-level balance sheets to construct events of abrupt changes in firm value-added, as in Lamadon et al. (2022). Such events capture both sudden firm-level changes in product demand and productivity over time, after purging common changes within the firm’s local labor market. Second, we use information on public procurement contracts, exploiting the fact that contract winners gain substantial revenue streams that in turn affect labor demand, as in Kroft et al. (2025) and Volpe (2024). Third, we use changes in international trade following the Great Recession to construct firm-level exposure to export-demand shocks, as in Garin and Silvério (2024) and Bhuller et al. (2025). We next discuss each of these research designs in more detail. We refer to the first design as using an “internal” demand shock, while the remaining two designs rely on demand shocks that are “external” to the firm.

Internal shock. Annual firm-level balance sheets provide information on firm revenues and value-added. We use the annual growth in value added from one year to the next to construct the internal shock, following Lamadon et al. (2022). Throughout, we index each event by its cohort c , the year when a firm is exposed to the shock. Let $\Delta VA_{cj} = \frac{VA_{cj,0} - VA_{cj,-1}}{VA_{cj,-1}}$ denote value-added growth for firm j in cohort c between event time -1 and 0 , and let ΔVA_c^{Median} denote the median value-added growth within the firm’s two-digit industry in the same cohort. We define the exposure to an internal demand shock as follows:

$$Z_{cj}^{Int} = \begin{cases} 1, & \text{if } \Delta VA_{cj} > \Delta VA_c^{Median}, \\ 0, & \text{if } \Delta VA_{cj} \leq \Delta VA_c^{Median}. \end{cases}$$

We implement the internal shock design on two different samples; Appendix Tables C.2–C.4 provide summary statistics for both of these samples. The first is our extended sample, comprising cohorts from 2005–2019, which includes 77,385 firms with complete information on exposure to the internal shock, along with linked employer–employee data, detailed firm cost accounts over a 7-year window, and a minimum of 2 stayers. The second one is our contracts sample, covering only the 2018–2019 cohorts, where there are 43,880 firms that can be linked to detailed matched employer–employee data with information on contract types and in-kind compensation, and for around half of these firms (21,114), we can also link information on mentions of workplace amenities from their posting of job ads.

Procurement shock. We use administrative records of all public procurement contracts awarded in Norway, from which we identify firms that win a government contract for the first time. Public procurement has been used as a source of variation in firm-level demand in Kroft et al. (2025) for U.S. construction firms. Volpe (2024) adapts this strategy to Norwegian firms, and we follow a similar approach. We restrict the sample to two-digit sectors in which at least five firms won their first procurement contract between 2012 and 2022, and use the universe of first-time winners in the 2012–2019 cohorts. We define exposure to the procurement shock as a firm j winning a government contract for the first time as:

$$Z_{cj}^{Proc} = \begin{cases} 1, & \text{if firm } j \text{ wins its first procurement contract in cohort } c, \\ 0, & \text{if firm } j \text{ never wins a procurement contract.} \end{cases}$$

The control group consists of never-treated firms within the same two-digit sectors, and was selected through a propensity score procedure.²⁵ Appendix Table C.2 provides summary statistics for the sample used for this shock. For the 2012–2019 cohorts, we use 18,243 firms in the procurement shock design, for which we can link matched employer–employee data with detailed firm cost accounts over a 7-year window, and have at least 2 stayers.

Export shock. We use changes in the patterns of international trade that took place following the Great Recession to construct a firm-level export demand shock exposure. This strategy follows Garin and Silvério (2024) for Portugal and Bhuller et al. (2025) for Norway. We link the matched employer–employee data to firm-level customs records that contain exported values by 8-digit HS product code and destination country. We combine this with publicly available bilateral trade flow data from the *BACI-CEPII* dataset (Gaulier and

²⁵We estimate a logit model for winning using data from one period before the firm wins, on pre-treatment employment, wages, value added, revenues, one-digit sector fixed effects, and commuting zone fixed effects, and restrict the control group to never-treated firms with a predicted winning probability above 1%.

Zignago, 2010), disaggregated at the HS6 product and country-pair level, to construct the country-product demand shifts that form the export shock design.

Following Garin and Silvério (2024) and Bhuller et al. (2025), we use a shift-share strategy that interacts firms' pre-period export exposure across country-product pairs with changes in international import demand around the 2007–2009 Great Recession, $Z_{cj}^{Ext} = \sum_{p \in \mathcal{P}} s_{jp} \Delta_p$, where the exposure shares $s_{jp} = \text{Exp}_{jp}^{2005-07} / \left(\sum_{p \in \mathcal{P}} \text{Exp}_{jp}^{2005-07} \right)$ are constructed from firm j 's real exported value to country-product p over 2005–2007. The symmetric growth rate is defined as $\Delta_p = (M_p^{\text{post}} - M_p^{\text{pre}}) / [(M_p^{\text{post}} + M_p^{\text{pre}}) / 2]$, where M_p denotes the total value of non-domestic imports (i.e., leaving out imports from Norway) of country-product p in the pre (2006–07) and post (2009–10) periods (Bhuller et al., 2025).

For the export shock, we focus on a single cohort that experienced changes in export demand around 2007–2009 and link information from matched employer–employee data and detailed firm cost accounts over a 10-year window from 2002 to 2011. Notably, the time periods and available information differ substantially from those in the remaining analysis, as contract-type indicators and job-ad data are not available for the period used in the export-shock design. Appendix Table C.2 provides summary statistics for the main sample of 3,957 firms used in the export shock design.²⁶

In Appendix Table C.3, we report the industry composition of each shock sample (including both treatment and control firms). The export shock is dominated by manufacturing (41.5%) and trade (37.6%), reflecting the concentration of exporting activity in goods-producing and wholesale sectors. The procurement shock sample is instead tilted toward construction (26.1%), trade (19.8%), and service sectors, consistent with the types of firms that typically win public contracts. The internal shock sample is more diversified, with trade (28.0%) and construction (17.5%) as the largest sectors. These differences in industry composition across designs should be kept in mind when comparing the shock responses.

3.3 Event Study Specification

For each shock, we estimate dynamic firm-level responses from event-study models as follows:

$$\ln Y_{jt} = \sum_{\tau \in \mathcal{T}} \beta_{Y,\tau} Z_{cj} \mathbf{1}\{\tau = t - c\} + \delta_{k(j)t}^Y + \gamma_j^Y + v_{jt}^Y, \quad (21)$$

where Y_{jt} represents different outcomes, including employment L_{jt} and mean wages w_{jt} , besides value added VA_{jt} and measures of amenities or non-wage compensation as discussed above, while $Z_{cj} \in \{Z_{cj}^{Int}, Z_{cj}^{Proc}, Z_{cj}^{Ext}\}$ denotes the relevant shock measure used in each shock

²⁶As in Bhuller et al. (2025), we restrict the sample to exporting firms whose exports accounted for at least 1% of revenues in 2005–2007, to capture firms in which exports constitute a meaningful economic activity.

design. The relative time-to-event year is indexed by τ , so $\mathbf{1}\{\tau = t - c\}$ is an indicator for event time relative to the shock year c , and $\mathcal{T}$ is the set of relative years included. The terms γ_j^Y and $\delta_{k(j)t}^Y$ denote firm fixed effects and local labor market by calendar year fixed effects, respectively, which control for constant firm-level parameters, such as baseline firm productivity, and isolate the effects from market-level demand shocks, which could otherwise confound the results. In practice, we define the local labor market $k(j)$ as the two-digit industry interacted with the commuting zone where each firm is located (Bhuller, 2009).²⁷

All event-study figures report coefficients with 95% confidence intervals. For the internal shock, the baseline period is the event time -2 , and we estimate the model using a balanced $\mathcal{T} = [-3, +3]$ event window, with standard errors clustered at the cohort and local labor market (LLM) level. For the procurement shock, we use the same event window, but the baseline period is event time -1 , and we cluster standard errors at the cohort and firm level. For these two designs, we restrict to firms observed in all event years that have at least two stayers. For the export shock, we use a single cohort window $\mathcal{T} = [2002, 2011]$, with a baseline period of 2006, and stayers are defined as workers present at the firm for at least 8 years within the estimation window, and we restrict to firms with at least two stayers. Standard errors in this design are clustered at the firm level. Lastly, as a summary statistic, we report the post-shock average effect as $\bar{\beta}_Y = (\hat{\beta}_{Y,1} + \hat{\beta}_{Y,2} + \hat{\beta}_{Y,3})/3$ for the internal and procurement shock, and as $\bar{\beta}_Y = (\beta_{Y,2010} + \beta_{Y,2011})/2$ for the export shock. For outcomes in logarithms, we report percentage changes, which we compute as $100 \times (e^{\bar{\beta}_Y} - 1)$.

Identification. The dynamic firm-level responses from the event-study models (21) provide reduced-form estimates of the impacts of demand shocks on employment, wages, and amenity measures. A causal interpretation of these coefficients requires additional identifying assumptions. Next, we provide conditions that support causal interpretation and discuss connections to our theoretical model in Section 2. We present these conditions in terms of a binary shock indicator Z_{cj} that applies directly to the internal and procurement shocks, whereas analogous conditions are needed for the continuous export shock.

Assumption 1 (Relevance). *Firms exposed to the shock experience differential changes in value added within their local labor market $k(j)$,*

$$\mathbb{E}[\Delta \ln V A_{jt} \mid Z_{cj} = 1, \delta_{k(j)t}^Y] \neq \mathbb{E}[\Delta \ln V A_{jt} \mid Z_{cj} = 0, \delta_{k(j)t}^Y],$$

where $Z_{cj} \in \{Z_{cj}^{Int}, Z_{cj}^{Proc}, Z_{cj}^{Ext}\}$. For the internal shock, this relevance condition is satisfied by construction. For the procurement and export shocks, it is an empirical requirement. In

²⁷For multi-establishment firms, we use the location of the establishment with most employees.

the context of our theoretical model in Section 2, this assumption implies that each shock Z_{cj} must be correlated with changes in the structural productivity parameter $\Delta\psi_{jt}$, which determines the location of the firm's labor demand curve.

Assumption 2 (Exogeneity). *Let $Y_{jt}(1)$ and $Y_{jt}(0)$ denote the potential outcomes with and without the shock, respectively. This assumption states that, absent the shock, the treated and control firms would have followed parallel trends, within their local labor market $k(j)$:*

$$\mathbb{E}[\Delta \ln Y_{jt}(0) \mid Z_{cj} = 1, \delta_{k(j)t}^Y] = \mathbb{E}[\Delta \ln Y_{jt}(0) \mid Z_{cj} = 0, \delta_{k(j)t}^Y].$$

This assumption states that the residual v_{jt}^Y is uncorrelated with the shock Z_{cj} . In the context of our theoretical model in Section 2, each shock Z_{cj} must be uncorrelated with changes in other firm-level determinants of wages and employment in (11) and amenity provision in (5), i.e., exogenous amenities $\Delta\bar{a}_{jt}$ and amenity costs $\Delta\mathbf{c}_{jt}$.

We provide evidence supporting Assumption 1, documenting statistically significant expansion of firms following each of the three demand shocks (Appendix Figure C.1). Averaged across the post-event periods, exposure to the internal demand shock increases firm value added by 29.6%, while a procurement win increases it by 10.7% and a one-unit increase in export demand exposure increases it by 34.0%. We also empirically assess parallel trends using pre-shock event-study coefficients, embedded in Assumption 2. Across outcomes and all three shock designs, we find no evidence of pre-trends, suggesting that the parallel trends assumption is plausible. We can thus interpret post-shock event-study estimates as capturing the causal reduced-form impacts of demand shocks on firm outcomes and amenity provision.

Meanwhile, as discussed in Section 2.5, the primary objects of interest in our analysis are the wage elasticity of the effective labor supply, ε_w^L , and the size elasticity of the amenity cost multiplier ε_L^A . Intuitively, the ratio of employment responses $\bar{\beta}_L$ to wage responses $\bar{\beta}_w$ provides a natural empirical analog for the effective labor supply elasticity, ε_w^L . This object can be equivalently estimated in an instrumental variables (IV) approach, where the demand shifter Z_{cj} serves as an instrument for the wage change $\Delta \ln w_{jt}$. Similarly, the ratio of amenity cost responses $\bar{\beta}_A$ to employment responses $\bar{\beta}_L$ provides a natural empirical analog for the amenity cost multiplier elasticity, ε_L^A . Assumption 2 further clarifies which amenity-related components may adjust without challenging the identification of ε_w^L and ε_L^A ; in an IV approach, the analogous conditions can be stated as exclusion restrictions.

Importantly, as elasticities can vary with firm size in our model, we quantify these objects in several steps. First, we split our estimation sample into ten parts by deciles of the pre-shock firm value-added distribution (and similarly by the firm size distribution). Second, we construct decile-specific reduced-form estimates of the employment and wage responses to

the demand shock. For this part of our analysis, we rely on the internal shock design. Next, using the decile-specific reduced-form estimates, we construct the ratios of employment-to-wage responses for each decile, which we interpret as measures of the underlying *local* labor supply elasticities $\varepsilon_w^L(L_j)$ across the firm value-added (or size) distribution. We follow a similar approach to quantify the amenity cost multiplier elasticities $\varepsilon_L^A(L_j)$ across the firm distribution. By combining the estimates of effective labor supply and amenity cost elasticities, we can finally recover decile-specific measures of the total compensation markdowns.²⁸ This combination is novel to our design, which uses employment, wage, and amenity responses to demand shocks across the firm distribution for model identification.

4 Evidence on Workplace Amenity Provision

We now provide causal evidence from three alternative event-study designs on adjustments in employment, wages and amenities in response to firm-level demand shocks.

4.1 Impacts on Wages and Employment


A key focus of the recent literature quantifying employer market power is how firm-level demand shocks affect employment and wages, which jointly inform the elasticity of firm-level labor supply in canonical monopsony models (Garin and Silvério, 2024; Lamadon et al., 2022). Before turning to our evidence on firm responses across various margins of non-wage compensation, we show the employment and wage responses induced by each shock, along with estimates of the implied firm-level labor supply elasticities.

Figure 2 reports employment and wage responses across the three demand shocks.²⁹ Averaged across the post-event period, we find that employment expands by 17.5% and stayers' wages increase by 4.5% using the internal shock (panels a–b). These reduced-form impacts provide an implied labor supply elasticity of $\varepsilon_w^L = \hat{\beta}_L / \hat{\beta}_w \approx 3.63$ and an implied wage pass-through elasticity of $\varepsilon_{VA}^w = \hat{\beta}_w / \hat{\beta}_{VA} \approx 0.171$. For the procurement shock (panels c–d), we find that employment expands by 8.2% and stayers' wages increase by 1.3%, while for the export shock (panels e–f), employment expands by 14.8% and stayers' wages increase by 7.7%. Appendix Table C.6 provides detailed point estimates across the three shocks.


²⁸Notably, our Proposition 1 clarifies that, unless $\rho = 0$ or $\Phi(L_j) = 0$, identifying the distribution of $\varepsilon_w^L(L_j)$ is not sufficient for the identification of the total compensation markdown $\mathcal{A}(L_j)\mathcal{M}(L_j)$. From (14), this requires identifying both the labor supply $\varepsilon_w^L(L_j)$ and the amenity cost multiplier $\varepsilon_L^A(L_j)$ elasticities.

²⁹As visualized across panels (a)–(c) of Appendix Figure C.1, all three shocks induce salient increases in firm value added, but the magnitudes differ. When interpreting the evidence on reduced-form impacts presented here, it is important to note these differences in relative shock intensities.


Figure 2: The Effects of Demand Shocks on Employment and Wages.




(a) Internal Shock: Employment




(b) Internal Shock: Wages




(c) Procurement Shock: Employment



(d) Procurement Shock: Wages



(e) Export Shock: Employment



(f) Export Shock: Wages

Notes: The main datasets used to construct outcomes are matched employer–employee data. Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Eq. (21). All outcomes are in logs. Wage outcomes are expressed in hourly terms and constructed from the stayers sample. Panels (a)–(b): internal shock for cohorts 2005–2019 with standard errors clustered at the LLM and cohort level. Panels (c)–(d): procurement shock for cohorts 2012–2019 with standard errors clustered at the firm and cohort level. For these two shocks, we restrict to firms observed in all years of the balanced $[-3, +3]$ event window and having at least two stayers. Panels (e)–(f): export shock with year 2006 as the baseline period and standard errors clustered at the firm level. Stayers in this shock are defined as workers who are present at the firm for at least 8 years during the estimation window. The export sample is restricted to firms with an export-to-revenue ratio of at least 1% in 2005–2007.

While the magnitudes of employment and wage impacts differ across the three shocks, partly reflecting differences in shock intensities (Appendix Figure C.1) and partly differences in firm composition (Appendix Table C.3), we find that firms respond significantly to each shock on both the employment and wage margins. The qualitative pattern is nevertheless consistent across all three designs: employment responds more than wages, and the implied elasticities are broadly similar despite differences in the sources of variation exploited.³⁰

4.2 Impacts on Non-Wage Compensation

Having established how firm-level demand shocks affect employment and wage margins, the remaining analysis examines whether firms also adjust the provision of workplace amenities in response to such shocks. In existing studies, this margin is typically assumed fixed or omitted, despite being an important component of the compensation bundle. We now use detailed information on categories of non-wage compensation to study this margin of response.

Figure 3 reports the responses of wage and non-wage components of compensation, measured using firms' cost accounts, to demand shocks. Panel (a) reports total personnel costs and panel (b) reports per-worker costs, both using the internal shock (contracts sample), highlighting heterogeneity across amenity components. Focusing on total personnel costs (panel a), we find that wage costs rise by 21.3%, pension premiums by 18.9%, and other non-wage personnel costs by 21.0%. Turning to per-worker personnel costs (panel b), we find that wage costs increase by 3.4%, pension premiums per worker by 1.7%, and other non-wage personnel costs by 3.1%. Panels (c)–(f) present the analogous decompositions of total and per-worker labor cost responses for the procurement shock and export shock designs, respectively. We again find that wage costs adjust across all three shocks with variation in the precision and size of the adjustments, both per worker and in total. Turning to total non-wage compensation for the procurement shock (panel c), we find qualitatively similar—yet quantitatively weaker—responses across all sub-components. Consistent with these results

³⁰In Appendix Figure C.3, panels (a)–(b), we provide employment and wage responses across the two internal-shock samples, namely the extended sample (2005–2019 cohorts) and the contracts sample (2018–2019 cohorts). We find very similar employment responses, with differences arising primarily in wage effects. Further, we re-estimate the internal shock specification for firms in the procurement-shock sample (panels (c)–(d)) and the export-shock sample (panels (e)–(f)), for comparison to the external shock estimates. Once sample differences are accounted for, we find similar employment and wage responses across the designs, suggesting that the different wage responses in the export- or procurement-shock sample reflect the composition of firms rather than the design of the shock itself. For the internal shock design (contracts sample), we provide additional evidence on employment and wage margins in Appendix Figure C.2, documenting that our findings are robust across alternative definitions of employment—considering the number of employees or measures of full-time equivalent employment—and we find very similar employment outcomes when we focus on all incumbent workers following their employment trajectories. Our main evidence on wage impacts focuses on stayers who remained employed with the same firm; however, in Appendix Figure C.2, we also show wage responses for hires and across all incumbent workers.

and the employment expansion, we find no changes in per-worker compensation across most categories of non-wage compensation. The evidence from the export shock in panels (e)–(f) similarly shows weaker—and less precise—responses for most components.

One interpretation for the heterogeneous responses across non-wage components, consistent with Proposition 2-(ii), is that per-worker spending responses vary with the cost elasticity $\varepsilon_L^{h_k}$. When wages and amenities enter as substitutes ($\rho > 0$) in worker preferences, components whose provision cost scales less than proportionally with employment exhibit positive per-worker responses, while those with sufficiently convex cost structures may decline even as total spending rises. As discussed in Section 2.5, to disentangle whether these differences reflect cost curvature or the structure of preferences, we need to observe shock responses at the level of amenity provision, which we return to below.


To summarize the non-wage personnel cost responses shown above, we now construct measures of the amenity cost multiplier $\mathcal{A}$ using the empirical counterparts of the objects in (6). Specifically, we use information on both non-wage compensation and wage bill from firms’ cost accounts. While firms may adjust workplace amenities that do not necessarily trigger changes in cost accounts, we view this measurement as an illustration of how information on amenity costs can guide the accounting of adjustments to such margins. Figure 4 shows the shock responses in the amenity cost multiplier $\mathcal{A}$. In panel (a), we find, using the internal demand shock, that the amenity cost multiplier, in percentage terms, adjusts by -0.12%, indicating that wages absorb a slightly larger fraction of the shock than non-wage costs. Using external shocks in panels (c) and (e), we also find declines in the multiplier by -0.15% for the procurement shock and -0.54% for the export shock, although with less precision. Panels (b), (d), and (f) decompose these responses by cost component, expressed as shares $\omega_k = C_k/WL$. Other non-wage personnel costs are the driver of these declines.

In Section 5, we return to how the impacts on amenity cost multipliers reported in Figure 4 guide our quantification of the firm distribution of total compensation markdowns, and in turn shape our evidence on employer market power with endogenous amenity responses.


4.3 Impacts on Contract Types and Advertised Amenities

Having documented adjustments in measures of non-wage compensation drawn from firm cost accounts, we now turn to worker-level information on in-kind compensation and contract types from matched employer–employee data, and then move on to study adjustments in workplace amenities advertised in job postings that we can match to our administrative data. As most of this information is available only for the latest years of our data collection, we focus solely on the 2018–2019 cohorts and exploit the internal shock design.


Figure 3: The Effects of Demand Shocks on Wage and Non-Wage Compensation.




(a) Internal Shock: Total Compensation




(b) Internal Shock: Per-Worker



(c) Procurement Shock: Total Compensation



(d) Procurement Shock: Per-Worker




(e) Export Shock: Total Compensation



(f) Export Shock: Per-Worker

Notes: The main datasets used to construct outcomes are firm cost accounts and matched employer–employee data (A-meldingen). Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Eq. (21). All outcomes are in logs. Panels (a), (c), and (e) report total compensation costs by type, while panels (b), (d), and (f) report per-worker compensation costs by type. In panels (a)–(b), the sample uses the 2018–2019 cohorts, is restricted to firms observed in all years of the balanced $[-3, +3]$ event window with at least two stayers, and standard errors are clustered at the LLM level. In panels (c)–(d), the sample uses the 2012–2019 cohorts, with the baseline period equal to -1 in the event study, and standard errors are clustered at the firm and cohort level. In panels (e)–(f), the sample includes firms with an export-to-revenue ratio of at least 1% in 2005–2007, with the baseline period set to 2006 and standard errors clustered at the firm level.


Figure 4: The Effects of Demand Shocks on the Amenity Cost Multiplier.




(a) Internal Shock




(b) Internal Shock: by Cost Component



(c) Procurement Shock



(d) Procurement Shock: by Cost Component




(e) Export Shock



(f) Export Shock: by Cost Component

Notes: The datasets used to construct outcomes are firm cost accounts. Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Eq. (21). The amenity cost multiplier $\mathcal{A} \equiv (\text{wage bill} + \text{total non-wage costs})/\text{wage bill}$ is defined as in Eq. (6). Panels (a), (c), and (e) report the overall amenity cost multiplier in levels, while panels (b), (d), and (f) decompose it by type of amenity, where each outcome is the share $\omega_k = C_k/WL$ in levels. The internal shock (panels a and b) uses the cohorts 2018–2019 with standard errors clustered at the LLM level. The procurement shock (panels c and d) uses cohorts 2012–2019, with standard errors clustered at the firm and cohort level. The export shock (panels e and f) standard errors are clustered at the firm level. We define stayers in this shock as workers who remain in the same firm for at least 8 years during the estimation window. The sample is restricted to firms with an export-to-revenue ratio of at least 1% in 2005–2007.

Figure 5: The Effects of Internal Shock on In-Kind Compensation.




Notes: The dataset used to construct outcomes is matched employer–employee data (A-meldingen). Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Eq. (21). In-kind compensation is measured in logs, while the extensive-margin outcome measures the share of stayers receiving in-kind compensation. The sample is restricted to firms observed in all years of the balanced $[-3, +3]$ event window, and with at least two stayers throughout the full estimation period. Estimates stack the 2018 and 2019 cohorts. Standard errors are clustered at the LLM level.

Figure 5 separates the in-kind compensation response into extensive and intensive margins among stayers. Panel (a) reports the share of stayers receiving in-kind compensation, with an increase of 0.8 pp. Panel (b) reports the log of in-kind pay conditional on receiving any, showing a 2.0% increase. Both margins respond positively to a positive shock, indicating that firms provide in-kind benefits to more workers and increase the benefit amounts.

Next, in panel (a) of Figure 6, we focus on contract types among stayers to document how firms adjust existing employment contracts following a demand shock. Full-time shares among stayers increase by 1.1 pp, while the shares of fixed (non-variable) hours contracts and regular daytime (non-shift) contracts remain largely unchanged. Focusing on the subsample of posting firms, panels (b)–(c) of Figure 6 examine workplace amenities offered in job postings, by measuring the shares of a firm’s posted job ads that mention each specific amenity type. Across amenity dimensions, mentions of job training increase by 2.4 pp and mentions of workplace quality increase by 2.2 pp following the shock, consistent with firms enhancing the non-wage bundle offered to prospective hires as firm-level demand rises. Next, we focus on employers offering workers the possibility to work flexible hours, which is arguably an example of a production-relevant amenity, where we find an adjustment of -1.7 pp.³¹ By comparison, mentions of inconvenient hours increase by 0.7 pp, whereas we find no significant responses in mentions of additional perks, financial benefits, career progression

³¹Specifically, in Appendix A.4, we offer a possible interpretation for the decline in flexible hours: a production-relevant amenity with a provision cost that is primarily forgone output rather than a monetary cost can decline after a positive shock, even if its monetary cost component is purely per-worker.

Figure 6: The Effects of Internal Shock on Contract Type and Advertised Amenities.



Notes: The datasets used to construct outcomes are matched employer–employee data (A-meldingen) and the job ads data (FINN/NAV), from 2015–2022. Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Equation (21). Outcomes in panel (a) measure the share of stayers in full-time (above 20 hours of work per week), fixed (non-variable) hours, and regular daytime (non-shift) contracts, while outcomes in panels (b)–(c) measure the shares of posted job ads with mentions of listed workplace amenities. The sample is restricted to firms observed in all years of the balanced $[-3, +3]$ event window, and with at least two stayers throughout the full estimation period. We further restrict to firms with job postings for the outcomes with job ads. Estimates stack the 2018 and 2019 cohorts. Standard errors are clustered at the LLM level.

opportunities, or work-from-home arrangements.

5 Implications for Employer Market Power

We now discuss the implications of our findings for measuring employer market power. The event-study estimates shown in Section 4 map into objects that are implied by our conceptual framework. For instance, using the reduced-form estimates of employment and wage responses measured across the firm distribution, we can recover the distribution of effective labor supply elasticities $\varepsilon_w^L(L_j)$. Importantly, our model predicts that $\varepsilon_w^L(L_j)$ may vary

across firms, depending on the cost-curvature index $\Phi(L_j)$, which ultimately depends on firm size. We next document heterogeneity in effective labor supply and amenity cost elasticities across the firm distribution, before turning to the quantification of employer market power.

Elasticities. To assess firm heterogeneity, we residualize pre-shock firm value added by two-digit industry and commuting zone, rank firms into ten deciles of baseline value added, and estimate internal-shock specifications separately for each decile to recover the implied decile-specific elasticities. Figure 7 shows that the effective wage-elasticity of labor supply $\varepsilon_w^L(L_j)$ rises steeply across the firm distribution (panel a), from about 2 at the bottom decile to about 7.5 at the top decile, suggesting that more productive firms face a more elastic effective labor supply. Focusing on the average labor supply elasticity estimate would thus mask this apparent heterogeneity. The size-elasticity of amenity cost multipliers $\varepsilon_L^A(L_j)$ across the firm distribution is also heterogeneous (panel b), being significantly negative in the bottom decile and statistically indistinguishable from zero in the rest of the firm distribution. Further, we repeat this exercise with firms ranked into ten deciles of residualized baseline employment (panels c-d), finding similar patterns across the firm distribution.


Structural parameters. The variation in $\varepsilon_L^A(L_j)$ reflects the joint influence of the cost-curvature index $\Phi(L_j)$ and the wage-amenity preference ρ . To identify them separately, we implement the structural estimation approach described in Appendix B.³² In particular, we use the relationship in (20) to recover the value of ρ and η from the decile-specific variation in ε_w^L and ε_L^A . Using a nonlinear least squares estimation procedure, we obtain $\hat{\rho} \approx 0.062$ with 95% CI [0.043, 0.085], indicating that wages and amenities are gross substitutes (though close to the Cobb–Douglas benchmark of $\rho = 0$), and $\hat{\eta} \approx 4.3$ with 95% CI [3.7, 5.3], indicating that the taste dispersion parameter is close to the average effective labor supply elasticity.

Markdowns. Using the structural estimates of $\hat{\rho}$ and $\hat{\eta}$, we can recover the implied distribution of firm-level cost-curvature and total compensation markdowns across the firm productivity distribution, as illustrated in Figure 8. We find that the cost-curvature index $\Phi(L_j)$ is negative over most of the distribution, implying a convex amenity cost structure and effective labor supply elasticities below the value of the structural parameter η . Mean-


³²Notably, Appendix B also characterizes what can be identified in the limiting case when $\varepsilon_L^A = 0$. Having $\varepsilon_L^A(L_j)$ close to zero in (13) could in our model reflect a cost structure that is close to per-worker on aggregate ($\Phi(L_j) \approx 0$) or preferences close to Cobb–Douglas ($\rho \approx 0$). More broadly, non-responsive amenity provision could also reflect lumpy adjustment. If changing an amenity involves a fixed cost, the firm leaves it at its pre-shock level (absorbing the shock in wages alone) unless the shock is large enough. Thus, a null amenity response may be more informative about the size of the shock than about the cost structure. As our estimates are averaged over several years with large shocks to firm value added, this concern is limited.

Figure 7: The Distribution of Effective Labor Supply and Amenity Cost Elasticities.

Ranked by Baseline Value Added




(a) Effective Labor Supply Elasticity (ε_w^L)




(b) Amenity Cost Elasticity (ε_L^A)

Ranked by Baseline Employment



(c) Effective Labor Supply Elasticity (ε_w^L)




(d) Amenity Cost Elasticity (ε_L^A)

Notes: Elasticities are constructed from the post-shock average effects $\bar{\beta}_Y$ following the dynamic specification in Eq. (21): the labor supply elasticity ε_w^L is the ratio of the employment and wage coefficients, and ε_L^A is the elasticity of the amenity multiplier with respect to employment. Confidence bands are 95% percentile intervals from a parametric bootstrap ($B = 2,000$) that draws each decile's coefficient vector from $N(\hat{\beta}^d, \hat{\Omega}_d)$ and recomputes the ratios; their asymmetry reflects the skewed sampling distribution of a coefficient ratio, most pronounced where the wage response is least precisely estimated. Panels (a) and (b) rank firms by residualized baseline log value added by two-digit industry and commuting zone. Panels (c) and (d) rank firms by baseline log employment, residualized in the same way. We estimate a separate dynamic specification within each decile. The sample covers internal shock cohorts 2005–2019.


while, in the top three deciles, $\Phi(L_j)$ turns positive, implying that amenity costs at the most productive firms exhibit scale economies. The implied total compensation markdowns rise from 64 percent of the marginal product in the bottom decile to 88 percent in the top decile, so the gap between total compensation and marginal product of labor is larger in less productive firms. Further, wage markdowns are 5.6 to 10.6 percentage points below the corresponding total compensation markdowns (Appendix Table B.1). The productivity

Figure 8: The Distribution of Total Compensation Markdowns and Cost Structure.

Ranked by Baseline Value Added




(a) Amenity Cost-Curvature Index ($\Phi(L_j)$)




(b) Total Compensation Markdown ($\mathcal{A}^*\mathcal{M}^*$)

Ranked by Baseline Employment



(c) Amenity Cost-Curvature Index ($\Phi(L_j)$)



(d) Total Compensation Markdown ($\mathcal{A}^*\mathcal{M}^*$)

Notes: Both objects are implied by the structural estimates $(\hat{\eta}, \hat{\rho})$ recovered in Appendix B (see Eq. (20)), evaluated decile by decile at the estimated effective labor supply elasticity ε_w^L . The total-compensation markdown is $\mathcal{A}^*\mathcal{M}^* = \eta/(\eta + 1 - \eta\Phi)$ and the cost-curvature index is $\Phi(L_j) = (1 - \rho)(1/\eta - 1/\varepsilon_w^L)$. Negative Φ indicates convex aggregate amenity costs. Confidence intervals are 95% parametric-bootstrap percentile bands. Panels (b) and (d) split the gap $1 - \mathcal{A}^*\mathcal{M}^*$ into the cost-curvature wedge and the conduct wedge of Section 2.4. Panels (a) and (b) rank firms by residualized baseline log value added by two-digit industry and commuting zone. Panels (c) and (d) rank firms by baseline log employment, residualized in the same way. We estimate a separate dynamic specification within each decile. The sample covers internal shock cohorts 2005–2019.

gradient is driven by the cost curvature rather than by the amenity cost multiplier, which varies far less across deciles ($\mathcal{A}$ lies between 1.08 and 1.14). In panels (c) and (d) of Figure 8, we repeat the decomposition with firms ranked by residualized baseline employment. The implied cost-curvature and markdown profiles are again very similar: the total-compensation markdown rises from 0.65 in the bottom employment decile to 0.89 in the top.

Decomposition. We further decompose the markdowns into cost-curvature and conduct wedges, as in (16). Panel (b) of Figure 8 stacks the two wedges, $1 - 1/(1 - \Phi(L_j))$ and $1/(1 - \Phi(L_j)) - \mathcal{A}^*\mathcal{M}^*$, at each firm’s observed employment. In the bottom decile, about two-thirds of the gap between compensation and marginal product reflects cost curvature.³³ At the top, the cost-curvature factor is above one, and conduct accounts for the whole gap and more. Thus, for the least productive firms, most of the gap between compensation and marginal product is attributable to the structure of amenity provision costs, not employer market power. The conduct factor lies between 0.79 and 0.85 across deciles, so conduct lowers compensation by 15 to 21 percent below what the cost structure alone would imply.

Discussion. Our evidence shows that effective labor supply elasticities increase and total compensation markdowns are compressed as we move along the firm distribution. Decomposing the markdowns, we find that employer conduct accounts for about a third of the gap at the bottom of the firm distribution and for all of the gap at the top. These conclusions rest on three modeling assumptions maintained in our analysis: (i) labor is untaxed, (ii) firms do not interact strategically, and (iii) workers are homogeneous. We now consider each assumption in turn. First, we can extend our baseline model with progressive wage taxation, leaving amenities untaxed (see details in Appendix A.3). With Norwegian tax progressivity, the wage markdown is unchanged, the markdown on compensation as workers value it falls by at most 1.8 percentage points, and the gradient of $\Phi(L_j)$ is virtually intact. Progressive labor taxation therefore leaves our interpretation unaltered.³⁴ Second, in the calibrated model of Berger et al. (2026), oligopsonistic competition accounts for little of the variation in wage markdowns across firms in the United States, so we expect its role to be limited here as well, although this aspect could be evaluated more closely in our context. Third, worker heterogeneity offers an alternative interpretation of our evidence. In related work, Volpe (2024) also finds that high-wage firms in Norway face more elastic labor supply, and attributes it to worker composition: wage-sensitive workers concentrate in high-paying firms, while workers valuing amenities relatively more concentrate at the bottom of the firm

³³The decomposition in panel (b) of Figure 8 uses the estimated values of $\hat{\eta} = 4.3$ and $\hat{\rho} = 0.062$. Appendix B.5 shows that under Cobb–Douglas preferences, where the level of $\Phi(L_j)$ is only set-identified, the requirement that no amenity has a cost elasticity below zero still bounds the share from below at 67%.

³⁴A related concern can be that the firm’s labor costs also include employer payroll taxes, which enter neither the wage bill nor our measure of non-wage compensation. As we show in Appendix A.3.1, our expression for the markdown holds for labor costs inclusive of payroll taxes, $(\mathcal{A}+t)w_j$: the total compensation markdowns are markdowns on that cost, and the wage markdowns in Table 1 and Appendix Table B.1 are the total compensation markdowns divided by the amenity multiplier plus t rather than by the multiplier alone, which at $t = 0.141$ lowers them to 0.52–0.69 across deciles.

distribution.³⁵ We provide an alternative, complementary explanation: the cost curvature of amenity provision varies across firms and is most convex at the bottom of the firm distribution.³⁶ Estimating labor supply elasticities within worker types could in principle allow separating worker sorting from cost curvature, since sorting predicts the same elasticity at every firm for a given type while cost curvature does not. We leave this for future work.

Implications. What are the implications of our findings for *conventional measures* of employer market power? To answer this, we abstract from firm heterogeneity and focus on the average effective labor supply elasticity, $\varepsilon_w^L = \bar{\beta}_L / \bar{\beta}_w$, and compare our estimates of total compensation markdowns with those implied by conventional measures. Using the internal shock, we find an average effective labor supply elasticity of $\varepsilon_w^L = 3.63$ for the extended sample (cohorts 2005–2019), while the contracts sample (cohorts 2018–2019) yields 4.68. By comparison, the procurement shock yields a larger estimate of 6.32, while the export shock yields a lower estimate of 1.86, likely reflecting the differences in firm composition across the shock designs. Using the expressions for total compensation and wage markdowns in Eqs. (14)–(15) of our model, we can construct alternative markdown measures over the range of labor supply elasticities reported above. To quantify the total compensation markdown, we also require an estimate of the average size-elasticity of the amenity cost multiplier, which, using the internal shock (extended sample), is $\varepsilon_L^A = \bar{\beta}_A / \bar{\beta}_L = -0.003$. Our event-study evidence provides the corresponding elasticity estimates using the external shocks. To quantify the wage markdown, we must further quantify the average amenity-cost multiplier, which, in our comprehensive internal-shock extended sample, is $\bar{A} \approx 1.09$.

Table 1 presents our quantification and decomposition of total compensation markdowns into components across the three shock designs and samples. Compared to the canonical benchmark, our total compensation markdown implies a correction of between 0.2 and 1.7 percentage points in our main samples (Columns 1, 4, and 6). This finding is consistent with non-wage compensation being close to neutral in terms of its cost structure feedback for the average firm, i.e., amenity provision costs on average scale nearly per-worker ($\Phi(L) \approx 0$).

³⁵In Berger et al. (2026), heterogeneity in the marginal value of wealth works the other way: richer workers are less wage-sensitive, so high-wage firms face less elastic supply (as found in the US and Brazil).

³⁶In a back-of-the-envelope calculation, we find that the cost-curvature wedge would matter even if all variation in labor supply elasticities across firms comes from other forces, such as worker sorting. Holding the cost-curvature index constant and letting η_d vary across deciles instead, the amenity spending moments require $\Phi = -0.058$, which would still account for between 15 percent of the gap between total compensation and marginal product at the bottom and 48 percent at the top. Notably, from (13), a constant cost-curvature index $\Phi(L_j) = \Phi$ implies a constant ε_L^A across the firm distribution, so its cross-decile mean of -0.004 , with $\hat{\rho} = 0.062$, gives $\Phi = (1 - \hat{\rho})(-0.004)/\hat{\rho} \approx -0.06$. The cost-curvature wedge is then $1 - 1/(1 - \Phi) \approx 0.054$ in every decile, with a total gap $1 - \mathcal{A}^* \mathcal{M}^*$ of 0.366 in the bottom decile and 0.114 in the top. In this calculation, we recomputed $\mathcal{A}^* \mathcal{M}^*$ for each decile η_d by solving for its supply elasticity, $1/\eta_d = 1/\varepsilon_w^L + \Phi/(1 - \rho)$.

Table 1: Comparisons to Conventional Measures of Employer Market Power.

	Internal Shock		Procurement Sample		Export Sample	
	Extended	Contracts	Internal	Procurement	Internal	Export
	Sample	Sample	Shock	Shock	Shock	Shock
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Estimates of Effective Elasticities</i>						
$\varepsilon_w^L = \bar{\beta}_L / \bar{\beta}_w$	3.63	4.68	6.09	6.32	3.02	1.86
$\varepsilon_{Aw}^L = \varepsilon_w^L / (1 + \varepsilon_w^L \varepsilon_L^A)$	3.67	4.85	5.87	7.19	3.09	2.01
$\varepsilon_L^A = \bar{\beta}_A / \bar{\beta}_L$	-0.003	-0.008	0.006	-0.019	-0.008	-0.039
$\varepsilon_{VA}^w = \bar{\beta}_w / \bar{\beta}_{VA}$	0.171	0.138	0.113	0.123	0.160	0.254
<i>Markdowns Under Endogenous Amenity Provision</i>						
Total Compensation $\mathcal{AM}$	0.786	0.829	0.854	0.878	0.756	0.667
Wage $\mathcal{M} = \mathcal{AM} / \bar{\mathcal{A}}$	0.718	0.758	0.765	0.786	0.657	0.580
$\mathcal{AM} - \mathcal{M}$ (pp)	6.8	7.1	8.9	9.1	9.9	8.7
<i>Comparisons to Conventional Markdowns</i>						
$\varepsilon_w^L / (\varepsilon_w^L + 1)$	0.784	0.824	0.859	0.863	0.751	0.651
Diff. Relative to $\mathcal{AM}$ (pp)	0.2	0.5	-0.4	1.4	0.4	1.7

Notes: This table uses Eqs. (14)-(15) for markdown calculation. Elasticities are constructed from the post-shock average effects $\bar{\beta}_Y$. The labor supply elasticity ε_w^L is the ratio of the employment and wage coefficients. ε_{Aw}^L is the elasticity of employment with respect to total compensation per worker. ε_L^A is the elasticity of the amenity multiplier with respect to employment. ε_{VA}^w is the wage pass-through elasticity, defined as the ratio of the wage and value-added coefficients. The total compensation markdown is $\mathcal{AM} = \varepsilon_w^L / [(\varepsilon_w^L + 1) + \varepsilon_w^L \varepsilon_L^A]$. The conventional benchmark omits the amenity correction: $\varepsilon_w^L / (\varepsilon_w^L + 1)$. The wage markdown is the total compensation markdown divided by the average amenity multiplier in the pre-treatment period ($\bar{\mathcal{A}}$) for each specific subsample. The extended sample refers to cohorts 2005–2019, while the contracts sample refers to cohorts 2018–2019. Columns (3)–(4) restrict to the procurement-shock sample, and columns (5)–(6) to the export-shock sample.

Our estimate of the cost-curvature index, $\Phi(L) \approx -0.05$, implied by the sample-average of ε_L^A , is virtually indistinguishable from zero in the middle part of the firm distribution. However, as seen in Figure 8, cost-curvature varies significantly across the firm distribution, so conventional measures that focus on the average firm would fail to capture this dispersion. Meanwhile, since firms’ spending on non-wage compensation implies a non-negligible amenity wedge ($\bar{\mathcal{A}} > 1$), the wage markdowns we recover are around 6.8 to 9.1 percentage points deeper than the total compensation markdowns, accounting for the main difference between our markdown measures and conventional markdown measures reported in the literature.

6 Conclusion

This paper has examined how firms adjust workplace amenities in response to demand shocks and what those adjustments imply for measuring employer market power. Our theoretical

framework embeds endogenous, multi-dimensional amenity provision into a tractable monopsony model in which the per-worker cost of each amenity can scale flexibly with firm size. From this structure, we derive a compact sufficient statistic—the amenity cost-curvature index—that summarizes how the amenity bundle shapes the firm’s effective labor supply elasticity and the total compensation markdown. The model clarifies that the level of amenity provision, per-worker amenity spending, and the amenity multiplier convey distinct information about preferences and cost structure, and that conventional wage-based markdowns differ from the welfare-relevant total compensation markdowns by a measurable wedge.

Bringing these predictions to Norwegian administrative data that combine matched employer–employee records, firm cost accounts, and the near-universe of job postings, we document that firms respond to positive demand shocks across essentially every margin of compensation we can measure. Three complementary identification strategies—abrupt within-firm value-added changes, first-time public procurement contract wins, and shift-share export demand shocks—yield broadly consistent patterns: employment responds more than wages, in-kind benefits and non-wage compensation rise both in total and per worker, the share of full-time contracts increases, and the content of job postings tilts toward training and workplace-quality amenities. At the same time, the amenity cost multiplier changes only modestly, consistent with cost-curvature that is close to per-worker on average.

The implications for the measurement of employer market power are twofold. First, the gap between total compensation and marginal product is not all conduct. The estimated total-compensation markdowns rise from roughly 64 to 88 percent of marginal product across the firm distribution, and decomposing them into a cost-curvature wedge due to the structure of provision costs and a conduct wedge due to employer market power shows that most of the gap at the least productive firms reflects the cost structure, while conduct lowers compensation by 15 to 21 percent relative to what that cost structure alone would imply. Second, the gap between wage markdowns and total compensation markdowns is substantial: our evidence shows that, across the firm distribution, wage markdowns are 5.6 to 10.6 percentage points below the corresponding total compensation markdowns. When markdowns are measured as the gap between wages and the marginal product (e.g., from a production function estimation), a meaningful fraction of what is attributed to employer rent extraction may be compensation workers receive in non-wage benefits. By comparison, the measurement challenge is less severe when markdowns are inferred from the inverse labor supply elasticity: our evidence shows that the correction arising from endogenous amenity feedback is small for the average firm, reassuring that studies using firm-level demand shocks to estimate elasticities are informative about employer market power even with endogenous amenity provision.

Several limitations in our work point to directions for further research. Our framework

treats firms as atomistic and workers as homogeneous in their preferences. Heterogeneity in tastes across workers—and the sorting it induces across firms with different amenity profiles—is a natural next step, with implications for the distributional incidence of changes in employer market power (Berger et al., 2026; Volpe, 2024). Moreover, our evidence focuses on private-sector firms in Norway, where the regulatory environment and collective bargaining shape the menu of available amenities (Arold et al., 2025; Lagos, forthcoming). Comparable evidence from other institutional settings would thus help establish the external validity of our findings. Our analysis also abstracts from informational frictions about amenities on either side of the market—whether firms accurately perceive workers’ valuations and how workers learn about non-wage job attributes over their careers—aspects that have begun to receive attention in recent work in this area (Adrjan et al., 2026; Andresen et al., 2026; Cordes and Müller, 2026; Sulka, 2025). We leave such extensions to future work.

References

- Adams-Prassl, A., Balgová, M., Qian, M., and Waters, T. (forthcoming). “Firm Concentration and Job Design: The Case of Schedule-Flexible Work Arrangements”. *Review of Economics and Statistics*.
- Adrjan, P., Balgová, M., Jäger, S., Jessen, J., and Sockin, J. (2026). *Job Ads as Signals: Evidence from a Priced Amenity and Worker Beliefs*. NBER Working Paper 35457.
- Andresen, M. E., Bhuller, M., and Løvgren, A. (2026). *Pay Beliefs and the Amenity-Pay Tradeoff*. Working Paper.
- Arold, B. W., Ash, E., MacLeod, W. B., and Naidu, S. (2025). *Worker Rights in Collective Bargaining*. NBER Working Paper 33605.
- Audoly, R., Bhuller, M., and Reiremo, T. A. (2024). *The Pay and Non-Pay Content of Job Ads*. Working Paper.
- Bagger, J. and Lentz, R. (2019). “An Empirical Model of Wage Dispersion with Sorting”. *Review of Economic Studies* 86 (1): 153–190.
- Berger, D. W., Herkenhoff, K. F., Jeong, J., and Mongey, S. (2026). *Labor Market Power with Worker and Firm Heterogeneity*. NBER Working Paper 35640.
- Bhuller, M. (2009). *Inndeling av Norge i arbeidsmarkedsregioner*. Notater 2009/24. Statistics Norway.
- Bhuller, M., Delgado-Prieto, L., Hermo, S., and Lorentzen, L. (2025). *Wage-Setting Constraints and Firm Responses to Demand Shocks*. Working Paper.
- Bils, M., Kaymak, B., and Wu, K.-J. (2025). *Robinson Meets Roy: Monopsony Power and Comparative Advantage*. NBER Working Paper 34165.

- Caldwell, S., Dube, A., and Naidu, S. (forthcoming). “Monopsony Makes It Big”. *Journal of Economic Literature*.
- Caldwell, S., Haegele, I., and Heining, J. (2025). *Firm Pay and Worker Search*. IAB-Discussion Paper 4/2025. Institute for Employment Research (IAB).
- Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). “Firms and Labor Market Inequality: Evidence and Some Theory”. *Journal of Labor Economics* 36 (S1): S13–S70.
- Charles, K. K., Johnson, M. S., Stephens Jr., M., and Lee, D. Q. (2022). “Demand Conditions and Worker Safety: Evidence from Price Shocks in Mining”. *Journal of Labor Economics* 40 (1): 47–94.
- Colonna, A., Pérez Pérez, J., and Rodrian, L. (2025). *Violation and Enforcement of Workplace Regulations: Evidence from Mexican Firm Inspections*. Working Paper.
- Cordes, S. and Müller, M. (2026). *Do Firms Know What Workers Want?* CRC TR 224 Discussion Paper. University of Bonn and University of Mannheim.
- Di Porto, E., Pagano, M., Pezone, V., Saggio, R., and Schivardi, F. (2026). *Careers and Wages in Family Firms: Evidence from Matched Employer-Employee Data*. Working Paper.
- Dube, A., Naidu, S., and Reich, A. D. (2022). *Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Walmart Workers*. NBER Working Paper 30441.
- Garin, A. and Silvério, F. (2024). “How Responsive Are Wages to Firm-Specific Changes in Labor Demand? Evidence from Idiosyncratic Export Demand Shocks”. *Review of Economic Studies* 91 (3): 1671–1710.
- Gaulier, G. and Zignago, S. (2010). *BACI: International Trade Database at the Product Level, 1994–2007*. Working Paper 2010-23. CEPII.
- Kline, P. (2025). “Labor Market Monopsony: Fundamentals and Frontiers”. *Handbook of Labor Economics* 6: 655–728.
- Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). “Who Profits from Patents? Rent-Sharing at Innovative Firms”. *Quarterly Journal of Economics* 134 (3): 1343–1404.
- Kroft, K., Luo, Y., Mogstad, M., and Setzler, B. (2025). “Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry”. *American Economic Review* 115 (9): 2926–2969.
- Lachowska, M., Mas, A., Saggio, R., and Woodbury, S. A. (2026). “Work Hours Mismatch”. *Econometrica* 94 (3): 991–1025.
- Lagos, L. (forthcoming). “Union Bargaining Power and the Amenity-Wage Tradeoff”. *American Economic Review*.
- Lamadon, T., Mogstad, M., and Setzler, B. (2022). “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market”. *American Economic Review* 112 (1): 169–212.

- Lee, J. M. and Taylor, L. O. (2019). “Randomized Safety Inspections and Risk Exposure on the Job: Quasi-Experimental Estimates of the Value of a Statistical Life”. *American Economic Journal: Economic Policy* 11 (4): 350–374.
- Maestas, N., Mullen, K. J., Powell, D., Wachter, T. von, and Wenger, J. B. (2023). “The Value of Working Conditions in the United States and Implications for the Structure of Wages”. *American Economic Review* 113 (7): 2007–2047.
- Manning, A. (2003). “The Real Thin Theory: Monopsony in Modern Labour Markets”. *Labour Economics* 10 (2): 105–131.
- Manning, A. (2021). “Monopsony in Labor Markets: A Review”. *ILR Review* 74 (1): 3–26.
- Mas, A. (2025). “Non-Wage Amenities”. *Handbook of Labor Economics* 6: 373–446.
- Mas, A. and Pallais, A. (2017). “Valuing Alternative Work Arrangements”. *American Economic Review* 107 (12): 3722–3759.
- Rosen, S. (1986). “The Theory of Equalizing Differences”. *Handbook of Labor Economics* 1: 641–692.
- Roussille, N. and Scuderi, B. (2025). *Bidding for Talent: A Test of Conduct in a High-Wage Labor Market*. NBER Working Paper 33848.
- Seegmiller, B. (2025). *Valuing Labor Market Power: The Role of Productivity Advantages*. Working Paper. Kellogg School of Management.
- Sockin, J. (2022). *Show Me the Amenity: Are Higher-Paying Firms Better All Around?* CESifo Working Paper 9842. CESifo.
- Sorkin, I. (2018). “Ranking Firms Using Revealed Preference”. *Quarterly Journal of Economics* 133 (3): 1331–1393.
- Sulka, T. (2025). *Providing Benefits to Uninformed Workers*. CRC TRR 190 Discussion Paper 566. Rationality and Competition.
- Volpe, O. (2024). *Job Preferences, Labor Market Power, and Inequality*. Tech. rep.
- Woodbury, S. A. (1983). “Substitution between Wage and Nonwage Benefits”. *American Economic Review* 73 (1): 166–182.
- Woodbury, S. A. and Huang, W.-J. (1991). *The Tax Treatment of Fringe Benefits*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Appendix

A Theoretical Appendix

This appendix provides derivations for the baseline model of Section 2 (Appendix A.1), compares the monopsonist's wage-amenity bundle to that of a value-taking firm (Appendix A.2), extends the model with progressive taxation (Appendix A.3) and production-relevant amenities (Appendix A.4), reviews existing models of imperfect competition with endogenous amenities (Appendix A.5), and provides numerical illustrations of our model (Appendix A.6).

A.1 Baseline Model: Mathematical Derivations

A.1.1 Labor Supply

Given the structure described in Section 2.1, we let idiosyncratic tastes be independent draws $\xi_{ij} \sim \text{Fréchet}(\eta, \bar{a}_j)$ with CDF $\exp(-\bar{a}_j \xi^{-\eta})$, and maintain that each worker selects the firm offering the highest utility. By standard discrete choice results (McFadden, 1974), the probability of choosing firm j is $P_j = \bar{a}_j V_j^\eta / \sum_{l \in \mathcal{J}} \bar{a}_l V_l^\eta$, and with N workers the law of large numbers gives

$$L_j = N \cdot P_j = \frac{N \bar{a}_j}{\mathbb{V}^\eta} V_j^\eta, \quad \mathbb{V}^\eta \equiv \sum_{l \in \mathcal{J}} \bar{a}_l V_l^\eta,$$

which is Equation (1). In the continuum limit $L_j/N \rightarrow 0$, firms treat $\mathbb{V}$ as fixed and the residual supply elasticity collapses from the exact $\eta(1 - L_j/N)$ to η .

A.1.2 Firm's Problem

Firm j is atomistic and takes $\bar{a}_j$ and $\mathbb{V}$ as given and chooses $(w_j, a_{j1}, \dots, a_{jn})$ to maximize

$$\pi_j = \psi_j L_j^\alpha - w_j L_j - \sum_k c_{jk} a_{jk} h_k(L_j) \quad \text{subject to} \quad L_j = (N \bar{a}_j / \mathbb{V}^\eta) V_j^\eta,$$

with $V_j = (w_j^\rho + \sum_k \gamma_k a_{jk}^\rho)^{1/\rho}$.

FOC for w_j . Differentiating profit at fixed $\mathbf{a}_j$, $\partial \pi_j / \partial w_j = 0$ gives

$$\text{MRP}(L_j) - w_j - \sum_m c_{jm} a_{jm} h'_m(L_j) = \frac{L_j}{\partial L_j / \partial w_j}. \quad (\text{A.1})$$

FOC for a_{jk} . Differentiating profit at fixed w_j and fixed a_{jm} , $m \neq k$, $\partial\pi_j/\partial a_{jk} = 0$ gives

$$\text{MRP}(L_j) - w_j - \sum_m c_{jm} a_{jm} h'_m(L_j) = \frac{c_{jk} h_k(L_j)}{\partial L_j / \partial a_{jk}}. \quad (\text{A.2})$$

MRS condition. Dividing (A.2) by (A.1) gives

$$\frac{\partial L_j / \partial a_{jk}}{\partial L_j / \partial w_j} = \frac{c_{jk} h_k(L_j)}{L_j}.$$

Since L_j depends on a_{jk} and w_j only through V_j , the left-hand side equals the ratio of the CES partials, $(\partial V_j / \partial a_{jk}) / (\partial V_j / \partial w_j) = \gamma_k (a_{jk} / w_j)^{\rho-1}$, yielding the MRS condition

$$\gamma_k \left(\frac{a_{jk}}{w_j} \right)^{\rho-1} = \frac{c_{jk} h_k(L_j)}{L_j}, \quad (\text{A.3})$$

matching (4). Rearranging gives $a_{jk}^* = w_j \kappa_k(L_j)$ with $\kappa_k(L_j) \equiv (\gamma_k L_j / (c_{jk} h_k(L_j)))^{1/(1-\rho)}$.

A.1.3 Derivation of the Wage Equation

We simplify the two non-trivial terms in (A.1). For the cost terms on the LHS, using $C_{jk} = c_{jk} a_{jk} h_k(L_j)$, $\varepsilon_L^{h_k} = L_j h'_k(L_j) / h_k(L_j)$, and $\omega_k(L_j) = C_{jk} / (w_j L_j)$:

$$\sum_k c_{jk} a_{jk} h'_k(L_j) = \sum_k C_{jk} \frac{\varepsilon_L^{h_k}}{L_j} = w_j \sum_k \omega_k(L_j) \varepsilon_L^{h_k}.$$

For the RHS, Equation (1) gives $\partial L_j / \partial V_j = \eta L_j / V_j$ and the CES aggregator at fixed $\{a_{jk}\}$ gives $\partial V_j / \partial w_j = (V_j / w_j)^{1-\rho}$. Combining via chain rule and using $V_j^\rho = w_j^\rho \mathcal{A}(L_j)$ from (8):

$$\frac{L_j}{\partial L_j / \partial w_j} = \frac{L_j}{(\eta L_j / V_j) (V_j / w_j)^{1-\rho}} = \frac{V_j^\rho w_j^{1-\rho}}{\eta} = \frac{\mathcal{A}(L_j) w_j}{\eta}.$$

Substituting both into (A.1), dividing by w_j , and using $\text{MRP}(L_j) = \alpha \psi_j L_j^{\alpha-1}$ yields the wage equation (9):

$$w_j = \frac{\alpha \psi_j L_j^{\alpha-1}}{1 + \mathcal{A}(L_j) / \eta + \sum_k \omega_k(L_j) \varepsilon_L^{h_k}}. \quad (\text{A.4})$$

A.1.4 Elasticities of Effective Labor Supply and Amenity Cost Multiplier

Log-differentiating the supply curve (1) gives $d \ln L_j = \eta d \ln V_j$. Log-differentiating (8) gives $d \ln V_j = (1/\rho) d \ln \mathcal{A}(L_j) + d \ln w_j$. Since $\mathcal{A}(L_j)$ depends on L_j alone, $d \ln \mathcal{A}(L_j) = \varepsilon_L^{\mathcal{A}} d \ln L_j$.

Combining:

$$d \ln L_j (1 - \eta \varepsilon_L^A / \rho) = \eta d \ln w_j,$$

yielding $\varepsilon_w^L = \eta / (1 - \eta \varepsilon_L^A / \rho)$. We now express ε_L^A explicitly. Log-differentiating $\omega_k(L_j) = \gamma_k \kappa_k(L_j)^\rho$ from (7) with $\kappa_k(L_j) \propto (L_j / h_k(L_j))^{1/(1-\rho)}$:

$$\frac{d \ln \omega_k(L_j)}{d \ln L_j} = \frac{\rho(1 - \varepsilon_L^{h_k})}{1 - \rho}. \quad (\text{A.5})$$

Since $\mathcal{A}(L_j) = 1 + \sum_k \omega_k(L_j)$, $d\mathcal{A}(L_j) = \sum_k d\omega_k(L_j)$ and $d \ln \mathcal{A}(L_j) = d\mathcal{A}(L_j) / \mathcal{A}(L_j) = \sum_k (\omega_k(L_j) / \mathcal{A}(L_j)) d \ln \omega_k(L_j) = \sum_k \sigma_k(L_j) d \ln \omega_k(L_j)$, recovering the multiplier elasticity (13),

$$\varepsilon_L^A = \sum_k \sigma_k(L_j) \frac{\rho(1 - \varepsilon_L^{h_k})}{1 - \rho} = \frac{\rho \Phi(L_j)}{1 - \rho}.$$

Hence $\varepsilon_w^L = \eta / (1 - \eta \Phi(L_j) / (1 - \rho))$.

The index $\Phi(L_j)$ itself drifts along supply. By the product rule,

$$\frac{d\Phi(L_j)}{d \ln L_j} = \sum_m (1 - \varepsilon_L^{h_m}) \frac{d\sigma_m}{d \ln L_j} - \sum_m \sigma_m \frac{d\varepsilon_L^{h_m}}{d \ln L_j}, \quad (\text{A.6})$$

where the first term equals $\frac{\rho}{1-\rho} \text{Var}_\sigma(1 - \varepsilon_L^{h_m})$: since $\sigma_m = \omega_m / \mathcal{A}(L_j)$, (A.5) and $\varepsilon_L^A = \frac{\rho}{1-\rho} \Phi(L_j)$ give $d\sigma_m / d \ln L_j = \sigma_m \frac{\rho}{1-\rho} [(1 - \varepsilon_L^{h_m}) - \Phi(L_j)]$, and summing against $(1 - \varepsilon_L^{h_m})$ gives $\frac{\rho}{1-\rho} [\sum_m \sigma_m (1 - \varepsilon_L^{h_m})^2 - \Phi(L_j)^2]$. Since $d \ln L_j = \varepsilon_w^L d \ln w_j$ along supply, differentiating $\varepsilon_w^L = \eta / (1 - \eta \Phi(L_j) / (1 - \rho))$ gives $d \ln \varepsilon_w^L / d \ln w_j = \eta^2 (1 - \rho) (d\Phi(L_j) / d \ln L_j) / (1 - \rho - \eta \Phi(L_j))^2$, which with (A.6) gives:

$$\left. \frac{d \ln \varepsilon_w^L}{d \ln w_j} \right|_{\text{supply}} = \frac{\eta^2}{(1 - \rho - \eta \Phi(L_j))^2} \left[\rho \text{Var}_\sigma(1 - \varepsilon_L^{h_k}) - (1 - \rho) \sum_k \sigma_k(L_j) \frac{d\varepsilon_L^{h_k}}{d \ln L_j} \right], \quad (\text{A.7})$$

where $\text{Var}_\sigma(1 - \varepsilon_L^{h_k}) \equiv \sum_k \sigma_k(L_j) (1 - \varepsilon_L^{h_k})^2 - \left(\sum_k \sigma_k(L_j) (1 - \varepsilon_L^{h_k}) \right)^2$. We refer to (A.7) as the *super-elasticity of effective labor supply* (Kline, 2025). The first term in the bracket is the *curvature-dispersion*, scaled by ρ and non-zero when cost-curvature elasticities differ across amenities; the second is the *curvature-drift*, the share-weighted drift of the $\varepsilon_L^{h_k}$ along the supply curve, scaled by $1 - \rho$ and non-zero whenever some h_k is non-iso-elastic.

A.1.5 Size-Elasticity of the Wage Markdown

Log-differentiating $w_j = \mathcal{M}(L_j) \alpha \psi_j L_j^{\alpha-1}$ yields the wage-equation elasticity

$$\left. \frac{d \ln w_j}{d \ln L_j} \right|_{\text{wage eq.}} = (\alpha - 1) + \varepsilon_L^{\mathcal{M}}.$$

It remains to compute $\varepsilon_L^{\mathcal{M}}$ explicitly. By (15), the wage markdown $\mathcal{M}(L_j) = w_j/\text{MRP}(L_j)$ is the inverse of the wage-equation denominator in (9), so

$$\varepsilon_L^{\mathcal{M}} \equiv \frac{d \ln \mathcal{M}(L_j)}{d \ln L_j} = -\mathcal{M}(L_j) \cdot \frac{d}{d \ln L_j} \left[1 + \frac{\mathcal{A}(L_j)}{\eta} + \sum_k \omega_k(L_j) \varepsilon_L^{h_k} \right].$$

Using $d\mathcal{A}/d \ln L_j = \sum_k d\omega_k/d \ln L_j$, the product rule on $\omega_k \varepsilon_L^{h_k}$, and substituting $d\omega_k/d \ln L_j = \omega_k \rho(1 - \varepsilon_L^{h_k})/(1 - \rho)$ from (A.5):

$$\varepsilon_L^{\mathcal{M}} = -\mathcal{M}(L_j) \left[\frac{\rho}{1 - \rho} \sum_k \omega_k(L_j) (1 - \varepsilon_L^{h_k}) \left(\frac{1}{\eta} + \varepsilon_L^{h_k} \right) + \sum_k \omega_k(L_j) \frac{d\varepsilon_L^{h_k}}{d \ln L_j} \right], \quad (\text{A.8})$$

which we refer to as the size-elasticity of the wage markdown. The first term is a *share-drift* channel: it vanishes at the per-worker benchmark ($\varepsilon_L^{h_k} = 1$ for all k) and at Cobb–Douglas preferences ($\rho = 0$). The second is a *curvature-drift* channel: it vanishes whenever individual cost technologies h_k are iso-elastic, regardless of ρ .

A.1.6 Conditions for Existence and Uniqueness

We impose two regularity assumptions. The cost functions h_k are twice continuously differentiable, with $\varepsilon_L^{h_k}$ bounded as $L_j \rightarrow 0$. Total compensation $\mathcal{A}(L_j)w_j(L_j)$ stays bounded as $L_j \rightarrow 0$ and bounded away from zero as $L_j \rightarrow \infty$.¹ We further maintain two conditions at every L_j , which deliver a unique interior equilibrium under the regularity assumptions.

- (i) Supply slopes up: $1 - \rho - \eta \Phi(L_j) > 0$. Automatic at $\Phi(L_j) \leq 0$; since $\varepsilon_L^{h_k} \geq 0$ implies $\Phi(L_j) \leq 1 - 1/\mathcal{A}(L_j)$, it is satisfied whenever the amenity share of total compensation is below $(1 - \rho)/\eta$.

¹For $\rho > 0$ both limit requirements are implied by (i) and (ii): compensation per worker has employment elasticity $1/\eta - \Phi(L_j) > \rho/\eta$ under (i), so it rises with firm size, and a diverging size elasticity of any h_k with a nonvanishing share would make $\varepsilon_L^{\mathcal{M}}$ diverge, contradicting (ii). For $\rho \leq 0$ they add mild limits at small firms: no h_k may have a size elasticity that diverges as $L_j \rightarrow 0$, which excludes only functions like e^{-1/L_j} , and under complements $\Phi(L_j)$ may in addition not exceed $1/\eta$ as $L_j \rightarrow 0$.

- (ii) Wage-equation slopes down: $\varepsilon_L^{\mathcal{M}} - (1 - \alpha) < 0$. Automatic at $\rho = 0$ with iso-elastic costs (where $\varepsilon_L^{\mathcal{M}} = 0$).

Along its labor supply curve the firm effectively chooses only L_j : inverting (10) gives the wage that recruits L_j workers, $w_j(L_j) = (L_j \mathbb{V}^\eta / N \bar{a}_j)^{1/\eta} \mathcal{A}(L_j)^{-1/\rho}$, the amenities follow from (5), and total compensation per worker is $\mathcal{A}(L_j)w_j(L_j)$ by (6). Profit is

$$\pi_j(L_j) = \psi_j L_j^\alpha - L_j \mathcal{A}(L_j) w_j(L_j). \quad (\text{A.9})$$

Revenue has derivative $\text{MRP}(L_j) = \alpha \psi_j L_j^{\alpha-1}$. Log-differentiating the compensation bill along supply gives

$$\frac{d \ln [L_j \mathcal{A}(L_j) w_j(L_j)]}{d \ln L_j} = 1 + \varepsilon_L^{\mathcal{A}} + \frac{1}{\varepsilon_w^L} = 1 + \frac{\rho \Phi(L_j)}{1 - \rho} + \frac{1}{\eta} - \frac{\Phi(L_j)}{1 - \rho} = 1 + \frac{1}{\eta} - \Phi(L_j),$$

by (13) and (12). The derivative of the compensation bill is this elasticity times compensation per worker, so the marginal cost of a hire is $\mathcal{A}(L_j)w_j(L_j)[1 + 1/\eta - \Phi(L_j)] = w_j(L_j)[1 + \mathcal{A}(L_j)/\eta + \sum_k \omega_k(L_j)\varepsilon_L^{h_k}] = w_j(L_j)/\mathcal{M}(L_j)$, the supply wage over the markdown (15). Hence

$$\pi'_j(L_j) = \text{MRP}(L_j) - \frac{w_j(L_j)}{\mathcal{M}(L_j)}. \quad (\text{A.10})$$

Profit rises where the marginal revenue product exceeds the marginal cost of a hire and falls where it falls short. Consider their ratio, $\mathcal{M}(L_j)\text{MRP}(L_j)/w_j(L_j)$, which equals one at the first-order condition $\pi'_j(L_j) = 0$. Its log-derivative is $(\alpha - 1) + \varepsilon_L^{\mathcal{M}} - 1/\varepsilon_w^L$, with $\varepsilon_L^{\mathcal{M}}$ from Appendix A.1.5 and $1/\varepsilon_w^L$ from (12), which regroups into the slope of the wage equation minus the slope of supply:

$$\frac{d \ln [\mathcal{M}(L_j)\text{MRP}(L_j)/w_j(L_j)]}{d \ln L_j} = \underbrace{\varepsilon_L^{\mathcal{M}} - (1 - \alpha)}_{\text{wage-equation slope}} - \underbrace{\frac{1 - \rho - \eta \Phi(L_j)}{\eta(1 - \rho)}}_{\text{supply slope}}. \quad (\text{A.11})$$

Under (ii) the first term is negative and under (i) the second is positive, so the log-ratio falls with employment. Two consequences follow. *Uniqueness*: the ratio $\mathcal{M}(L_j)\text{MRP}(L_j)/w_j(L_j)$ is strictly decreasing, so it equals one at most once. *Existence*: by the regularity assumptions and $\Phi(L_j) < 1$, the marginal cost of a hire is bounded at small L_j and bounded away from zero at large L_j , while $\alpha < 1$ sends $\text{MRP}(L_j)$ to infinity and to zero at the same ends. The ratio thus starts above one, ends below it, and by continuity equals one at an interior L_j^* . Writing (A.10) as $\pi'_j(L_j) = \frac{w_j(L_j)}{\mathcal{M}(L_j)}[\mathcal{M}(L_j)\text{MRP}(L_j)/w_j(L_j) - 1]$, the bracket vanishes at L_j^* ,

so only its derivative enters $\pi_j''(L_j^*)$, and by (A.11) the second-order condition holds,

$$\pi_j''(L_j^*) = \frac{w_j(L_j^*)}{L_j^* \mathcal{M}(L_j^*)} \left[\varepsilon_L^{\mathcal{M}} - (1 - \alpha) - \frac{1 - \rho - \eta \Phi(L_j^*)}{\eta(1 - \rho)} \right] < 0.$$

Profit therefore rises on $(0, L_j^*)$ and falls beyond, so L_j^* is the global and interior optimum, and since the compensation bill vanishes with L_j , $\pi_j(L_j^*) > 0$: the firm prefers operating to shutdown.

Conditions (i) and (ii) are sufficient rather than necessary. The argument uses only that the log-ratio in (A.11) falls with employment, and (i) and (ii) sign its two pieces separately. The unique interior equilibrium therefore survives even where the wage falls with firm size, as long as the ratio still falls. We impose (i) for simplicity, since the comparative statics divide by $d \ln w_j > 0$.

Value-taking benchmark. Proposition A.1 in Section A.2 below compares the monopsonist with a firm that takes the job value $\bar{V}_j$ as given. The comparison requires that this firm also has a unique interior optimum. Its compensation bill is $L_j \mathcal{A}(L_j) w_j$ with w_j set so that the job delivers $\bar{V}_j$. By (13), the bill per worker falls with employment at rate $\Phi(L_j)$, and the marginal cost is (A.12). Profit is single-peaked, with an interior maximum, when this marginal cost falls more slowly with employment than the marginal revenue product, uniformly in L_j and by a strictly positive margin:

$$(ii') \quad \Phi(L_j) + \frac{1}{1 - \Phi(L_j)} \frac{d\Phi(L_j)}{d \ln L_j} < 1 - \alpha.$$

Condition (ii') is the counterpart of (i) and (ii) for a firm that does not internalize the supply term $\mathcal{A}(L_j) w_j / \eta$ in (A.13). Without that term, only diminishing returns in production work against scale economies in providing amenities, and (ii') requires that they prevail. We maintain (ii') wherever Proposition A.1 is invoked.

A.2 Comparison to the Competitive Benchmark

To see how employer market power shapes the composition of the wage-amenity bundle, consider the benchmark of *value-taking* conduct: the firm takes the job value $\bar{V}_j$ it must deliver as given and chooses employment L_j and pays the wage $w_j = \bar{V}_j \mathcal{A}(L_j)^{-1/\rho}$. The optimality conditions (4) remain unchanged, so the two conducts offer the same composition at the same scale and differ only in position on a common expansion path. By (8), delivering $\bar{V}_j$ costs $\mathcal{A}(L_j) w_j = \bar{V}_j \mathcal{A}(L_j)^{(\rho-1)/\rho}$ per worker, which with $\bar{V}_j$ fixed varies with employment

only through $\mathcal{A}(L_j)$. Differentiating the bill $L_j \cdot \bar{V}_j \mathcal{A}(L_j)^{(\rho-1)/\rho}$ by the product rule and using (13) gives the value-taker's marginal cost of hiring an additional worker:

$$\text{MC}(L_j) = \mathcal{A}(L_j) w_j \left[1 + \frac{\rho-1}{\rho} \varepsilon_L^{\mathcal{A}} \right] = [1 - \Phi(L_j)] \mathcal{A}(L_j) w_j. \quad (\text{A.12})$$

The value-taker's marginal cost can be decomposed into two terms: the first term $\mathcal{A}(L_j) w_j$ is the marginal hire's compensation, while the second term $-\Phi(L_j) \mathcal{A}(L_j) w_j$ captures the change in provision costs for all inframarginal workers as size increases, which lowers marginal cost under concave costs and raises it under convex costs. The value taker expands until $\text{MRP}(L_j) = \text{MC}(L_j)$, with employment L_j^c , wage w_j^c , and amenities $a_{jk}^c = a_{jk}(L_j^c)$.

While the monopsonist similarly expands until $\text{MRP}(L_j) = \text{MC}(L_j)$, its marginal cost has a different structure. Intuitively, the monopsonist cannot hire at a fixed value, since expanding means climbing its supply curve, so the marginal hire raises the value for every inframarginal worker. Thus, the value a monopsonist delivers is no longer fixed but rises along the supply curve, $d \ln V_j / d \ln L_j = 1/\eta$. Differentiation by the product rule adds to (A.12) an additional term capturing this rising value, $\mathcal{A}(L_j) w_j / \eta$:

$$\text{MC}(L_j) = [1 - \Phi(L_j)] \mathcal{A}(L_j) w_j + \frac{\mathcal{A}(L_j) w_j}{\eta}. \quad (\text{A.13})$$

As this term is positive, it lifts marginal cost above the value-taker's at every scale, so at L_j^c the monopsonist's profit is already falling. With a unique optimum for each conduct (Appendix A.1.6), it stops strictly short, $L_j^* < L_j^c$. At this smaller scale, the wage is lower ($w_j^* < w_j^c$), the job value is also lower ($V_j \propto L_j^{1/\eta}$), and the marginal revenue product higher ($\text{MRP}(L_j^*) > \text{MRP}(L_j^c)$). However, the provision of amenities can move either way. Both firms lie on the common path $a_{jk} = w_j (\gamma_k L_j / (c_{jk} h_k(L_j)))^{1/(1-\rho)}$ from (5); log-differentiating it along the supply curve, where $d \ln w_j / d \ln L_j = 1/\varepsilon_w^L > 0$, we have:


$$\left. \frac{d \ln a_{jk}}{d \ln L_j} \right|_{\text{supply}} = \frac{1}{\varepsilon_w^L} + \frac{1}{1-\rho} (1 - \varepsilon_L^{h_k}). \quad (\text{A.14})$$

Writing $a_{jk}^* = a_{jk}(L_j^*)$ and $a_{jk}^c = a_{jk}(L_j^c)$, the gradient is positive and the monopsonist under-provides the amenity ($a_{jk}^* < a_{jk}^c$) unless provision costs are strongly convex, $\varepsilon_L^{h_k} - 1 > (1-\rho)/\varepsilon_w^L$, in which case it *over*-provides it ($a_{jk}^* > a_{jk}^c$). Proposition A.1 summarizes.

Proposition A.1 (Provision of Wages and Amenities under Monopsony). *Let L_j^c be the optimal employment and let $(w_j^c, \{a_{jk}^c\})$ be the optimal wage-amenity bundle offered under a value-taking conduct. (i) Monopsony strictly under-provides employment, wages, and hence job value ($L_j^* < L_j^c$, $w_j^* < w_j^c$, $\text{MRP}(L_j^*) > \text{MRP}(L_j^c)$). (ii) Both conducts satisfy the same*

MRS condition (4) along an expansion path, so equilibrium amenities differ only due to the scale difference ($L_j^* < L_j^c$). **(iii)** Amenity k is under-provided, $a_{jk}^* < a_{jk}^c$, if it rises with scale along the expansion path, $\varepsilon_L^{h_k} - 1 < (1 - \rho)/\varepsilon_w^L$ throughout $[L_j^*, L_j^c]$ (per-worker, concave, or mildly convex costs), and over-provided if it falls with scale throughout that interval (strongly convex costs).

Appendix Figure A.1: Wage–Amenity Determination under Monopsony.



Notes: Wage–amenity packages for a single amenity k (with several amenities the figure applies amenity by amenity, since (4) holds for each k). Blue curves are iso-value loci of $V_j = (w_j^\rho + \gamma_k a_{jk}^\rho)^{1/\rho}$ at $\rho = 0.1$, $\gamma_k = 1$. The two marked bundles are the monopsony optimum (a_{jk}^*, w_j^*) , delivering job value V_j^* , and the value-taking optimum (a_{jk}^c, w_j^c) , delivering the higher V_j^c . Red iso-cost lines have slope $-c_{jk}h_k(L_j)/L_j$, solid at the monopsony bundle and dashed at the competitive one. The green expansion path traces the tangencies as L_j rises along the supply curve. In (d) the per-worker cost rises steeply enough that the gradient (A.14) turns negative at large L_j , bending the path backward.

Figure A.1 illustrates Proposition A.1 in the (w_j, a_{jk}) plane. Both bundles lie on the

common expansion path, with the monopsony optimum at the lower job value (parts (i) and (ii)), and cost curvature governs how the composition shifts along it: at per-worker costs the path is a straight ray and monopsony a radial contraction (panel a), concave costs tilt it toward the amenity (panel b), and convex costs tilt it toward wages (panels c–d). Under mild convexity the amenity still rises with scale (panel c), but once the gradient (A.14) turns negative the path bends backward (panel d), the over-provision case of part (iii).

A.3 Model Extension: Progressive Taxation

Our baseline model treats wages and amenities symmetrically. In this extension we tax wages by the log-linear schedule of Heathcote et al. (2017), $y - T(y) = \lambda y^{1-\tau}$, with progressivity $\tau \in [0, 1)$ and level $\lambda > 0$, leaving amenities untaxed and firm costs unchanged. Workers then value the disposable wage $g(w_j) \equiv \lambda w_j^{1-\tau}$, so the job-value index becomes $V_j = (g(w_j)^\rho + \sum_k \gamma_k a_{jk}^\rho)^{1/\rho}$. Repeating the FOC steps of Appendix A.1, the tax enters only through the rate at which gross wages pass into worker value. Progressivity makes that rate differ at the margin, $g'(w_j) = (1 - \tau)\lambda w_j^{-\tau}$, from the average, $g(w_j)/w_j = g'(w_j)/(1 - \tau)$, and both appear in the amenity–wage ratio:

$$\frac{a_{jk}^*}{w_j^*} = \frac{\Theta}{1 - \tau} \left[\left(\frac{\gamma_k}{c_{jk}} \right) \cdot \left(\frac{L_j}{h_k(L_j)} \right) \right]^{1/(1-\rho)}, \quad \Theta \equiv [g'(w_j^*)]^{-\rho/(1-\rho)}, \quad \omega_k^t = \frac{\Theta}{1 - \tau} \omega_k(L_j),$$

with $\omega_k(L_j)$ the baseline share from (7). The tilt Θ is the substitution response to the marginal rate, which any tax moves; the wedge $1/(1 - \tau)$ is the gap between average and marginal rates, which opens only under progressivity.

At $\tau = 0$, the wedge is one, and the tilt collapses to the constant $\Theta = \lambda^{-\rho/(1-\rho)}$: under substitutes, a proportional tax shifts the package toward the untaxed amenities, but by the same proportion at every firm. The model is then the baseline with rescaled shares $\omega_k^t = \Theta \omega_k$ —the wage FOC, both elasticities, the markdown (14) and the thresholds of Proposition 2 all keep their form, evaluated at curvature index $\Phi^t = \sum_k \sigma_k^t (1 - \varepsilon_L^{h_k})$ and shares $\sigma_k^t = \Theta \omega_k(L_j)/\mathcal{A}^t$. For $\tau > 0$, by contrast, the tilt depends on the equilibrium wage, $\Theta = [(1 - \tau)\lambda]^{-\rho/(1-\rho)} (w_j^*(L_j))^{\tau\rho/(1-\rho)}$ with $d \ln \Theta / d \ln w_j^* = \tau\rho/(1 - \rho)$, so under substitutes the tax channel tilts larger firms further toward the untaxed amenities, a cross-firm gradient in the amenity share that is absent from the baseline at per-worker costs and reversed under complements. The remaining objects inherit the tilt: the multiplier becomes $\mathcal{A}^t = 1 + \Theta \sum_k \omega_k(L_j)$, with value shares σ_k^t and cost curvature index Φ^t . Note that $\mathcal{A}^t$ measures what workers receive, not what the firm spends, and for $\tau > 0$ the two differ because the marginal wage is taxed more heavily than the average: amenity spending is

chosen against wage at the marginal rate $g'(w_j)$, while each wage unit delivers value at the average rate $g(w_j)/w_j = g'(w_j)/(1 - \tau)$, so spending per wage unit overstates the value contribution by $1/(1 - \tau)$. Combining ω_k^t with $\mathcal{A}^t$ gives $\mathcal{A}^t = 1 + (1 - \tau)(\hat{\mathcal{A}} - 1)$, where $\hat{\mathcal{A}} \equiv 1 + \sum_k \omega_k^t$ is the multiplier as measured in expenditure data and exceeds $\mathcal{A}^t$.

The firm's problem goes through as in the baseline but the value index collapses in the disposable wage, $V_j = g(w_j^*)(\mathcal{A}^t)^{1/\rho}$, and the wage FOC becomes

$$w_j = \frac{\alpha \psi_j L_j^{\alpha-1}}{1 + \mathcal{A}^t / [\eta(1 - \tau)] + \sum_k \omega_k^t \varepsilon_L^{h_k}},$$

where $1/(1 - \tau)$ reflects the average-marginal wedge. Log-differentiating supply, using $d \ln g(w_j) = (1 - \tau) d \ln w_j$ and $d \ln \Theta = \frac{\tau \rho}{1 - \rho} d \ln w_j$, gives the following two elasticities:

$$\varepsilon_w^L = \frac{\eta \left(1 - \tau \frac{1/\mathcal{A}^t - \rho}{1 - \rho}\right)}{1 - \eta \Phi^t / (1 - \rho)}, \quad \varepsilon_L^{\hat{\mathcal{A}}} = \rho \frac{\mathcal{A}^t}{\hat{\mathcal{A}}} \left[\frac{1}{\eta(1 - \tau)} - \frac{1}{\varepsilon_w^L} \right].$$

Progressivity scales the supply elasticity by the numerator factor $1 - \tau(1/\mathcal{A}^t - \rho)/(1 - \rho)$, below one whenever $\rho < 1/\mathcal{A}^t$, as at our estimates: raising the gross wage is a blunter instrument when part of it is taxed away. The second expression eliminates Φ^t and is stated for the multiplier as measured in expenditure data. Both reduce to their baseline forms at $\tau = 0$.

Proposition 1' (Markdowns under Progressive Wage Taxation). *Let wages be taxed by $y - T(y) = \lambda y^{1-\tau}$ with amenities untaxed. (i) The value-weighted markdown keeps the form of (14) with an adjusted curvature index and a numerator wedge,*

$$\mathcal{A}^t \mathcal{M} = \frac{\eta(1 - \tau)}{\eta + 1 - \eta \tilde{\Phi}^t}, \quad \tilde{\Phi}^t \equiv \Phi^t + \tau/\mathcal{A}^t, \quad (\text{A.15})$$

(ii) The wage markdown stays below it by the factor $\mathcal{A}^t > 1$, which for $\tau > 0$ lies below the expenditure multiplier $\hat{\mathcal{A}}$. (iii) Under $\tau > 0$ the canonical formula no longer recovers the markdown, even at $\rho = 0$. All parts reduce to Proposition 1 at $\tau = 0$, $\lambda = 1$.

Substituting the supply elasticity into (A.15) gives the total compensation markdown implied by a measured elasticity,

$$\mathcal{A}^t \mathcal{M} = \frac{1 - \tau}{1 + \rho/\eta - \tau/\mathcal{A}^t + [(1 - \rho) - \tau(1/\mathcal{A}^t - \rho)]/\varepsilon_w^L}, \quad \mathcal{A}^t = 1 + (1 - \tau)(\hat{\mathcal{A}} - 1),$$

so the right side involves only the schedule's τ , the measured elasticity, and the expenditure

multiplier $\hat{\mathcal{A}}$. Even at $\rho = 0$ a correction remains: the formula reduces to $\varepsilon_w^L/(1 + \varepsilon_w^L)$ scaled by $\mathcal{A}^t/\hat{\mathcal{A}}$.

Proposition 2' (Productivity Shocks under Progressive Wage Taxation). *Holding (τ, λ) fixed, after a positive productivity shock $d\psi_j > 0$: (i) The amenity level a_{jk}^* rises iff $\varepsilon_L^{h_k} - 1 < [1 - \rho(1 - \tau)]/\varepsilon_w^L$. (ii) Per-worker spending C_{jk}^*/L_j^* rises iff $\rho(\varepsilon_L^{h_k} - 1) < [1 - \rho(1 - \tau)]/\varepsilon_w^L$. (iii) The multiplier $\mathcal{A}^{t*}$ rises iff $\rho[\Phi^t + \tau(1 - 1/\mathcal{A}^t)/\varepsilon_w^L] > 0$, and is invariant at $\rho = 0$. All parts reduce to Proposition 2 at $\tau = 0$.*

Relative to the baseline thresholds, progressivity enters twice: $1 - \rho$ is replaced by $1 - \rho(1 - \tau)$ in parts (i)–(ii), and ε_w^L takes its tax-shifted value. Because the tilt rises with the equilibrium wage and the shock's own wage response pushes the package toward the untaxed amenities, loosening every threshold under substitutes and tightening it under complements. Since amenity costs do not depend on output here, the underprovision margin of Proposition A.1(iii) inherits the threshold of part (i).

A.3.1 Application: Structural Estimation with Tax Progressivity

How much do the tax progressivity corrections matter for our findings? Rearranging the expression for $\varepsilon_L^{\hat{\mathcal{A}}}$ above gives the functional relationship (20) with tax-adjusted coefficients,

$$\frac{1}{\varepsilon_w^L} = \frac{1}{\eta(1 - \tau)} - \frac{1}{\rho} \frac{\hat{\mathcal{A}}}{\mathcal{A}^t} \varepsilon_L^{\hat{\mathcal{A}}},$$

so the estimator runs unchanged and only its outputs relabel: $\eta = \hat{\eta}/(1 - \tau)$, while $\rho = \hat{\rho}\hat{\mathcal{A}}/\mathcal{A}^t$, evaluated at the cross-decile mean of the ratio since $\mathcal{A}_d^t$ varies slightly across deciles. We calibrate $\tau = 0.054$ for Norwegian income taxes (Qiu and Russo, 2026) and $\tau = 0.169$ once transfers are included (Holter et al., 2019), and compute $\mathcal{A}^t$ decile by decile from observed expenditure. As we treat all non-wage compensation as untaxed, which gives the exemption of amenities its largest possible effect, this correction is best read as an upper bound on the role of that exemption; some in-kind benefits are in fact taxable. Appendix Table A.1 shows the results: the markdown on compensation as workers value it, $\mathcal{A}^t\mathcal{M}$, falls by at most 1.8 percentage points, which is the revaluation of amenity spending at the net-of-tax rate; the expenditure-based markdown $\hat{\mathcal{A}}\mathcal{M}$ is unchanged, and its decile gradient is intact; the wage markdown is invariant to τ : substituting the relabeled parameters, every τ cancels and $\mathcal{M} = [\hat{\mathcal{A}}(1 + \hat{\rho}/\hat{\eta} + (1 - \hat{\rho})/\varepsilon_w^L)]^{-1}$, a function of measured objects alone. Re-inverting the decile-level curvature index under the two schedules moves Φ toward zero at every decile, by at most 0.02 at $\tau = 0.054$ and 0.05 at $\tau = 0.169$, with no sign change,

Appendix Table A.1: Structural Estimates under Calibrated Tax Progressivity.

	ρ	η	Φ^t (d1–d10)	$\mathcal{A}^t\mathcal{M}$ (d1–d10)	$\mathcal{M}$ (d1–d10)
$\tau = 0$	0.062	4.28	−0.325–0.095	0.641–0.878	0.585–0.772
$\tau = 0.054$	0.062	4.52	−0.310–0.089	0.638–0.872	0.585–0.772
$\tau = 0.169$	0.063	5.15	−0.278–0.077	0.632–0.860	0.585–0.772
$t = 0.141$	0.062	4.28	−0.325–0.095	0.641–0.878	0.518–0.687

Notes: Baseline estimates re-interpreted under the calibrated schedules: $\tau = 0.054$ is tax-only progressivity (Qiu and Russo, 2026), $\tau = 0.169$ adds transfers (Holter et al., 2019), both for Norway. The last row keeps $\tau = 0$ and applies the employer payroll tax rate $t = 0.141$ to the wage markdown. Ranges span productivity deciles 1–10 (value-added ranking).

so neither the convexity at the bottom nor the scale economies at the top are artifacts of untaxed amenities.

Employer payroll taxes. Section 3 excludes the employer payroll tax from non-wage compensation. However, a proportional payroll tax at rate t is equivalent to a proportional tax on wages, since only the ratio of the firm’s wage cost to the worker’s wage enters the firm’s problem and the supply curve: it is the case $\tau = 0$, $\lambda = 1/(1+t)$ above, which leaves ρ , η , $\Phi(L_j)$ and the total-compensation markdown unchanged. It does affect the interpretation of the markdowns, because the firm accounts record the wage before the employer’s contribution. The firm’s labor cost per worker is $(\mathcal{A} + t)w_j$, so the total-compensation markdown of Section 5 is a markdown on labor cost inclusive of the payroll tax, and the wage markdown is the total-compensation markdown divided by $\mathcal{A} + t$ rather than by $\mathcal{A}$ alone. The last row of Appendix Table A.1 applies the standard rate $t = 0.141$: the wage markdowns fall by about a tenth, from 0.59–0.77 to 0.52–0.69 across deciles.

Across the rows of Appendix Table A.1 the measured elasticities are held fixed; only the schedule assumed when interpreting them changes, so the $\tau = 0$ row is what one would conclude by ignoring taxation, not a no-tax economy. The effect of the tax itself comes from comparing two otherwise identical economies with and without the schedule: at $\tau = 0.169$ the supply elasticity is lower by 16 percent, through the factor $1 - \tau(1/\mathcal{A}^t - \rho)/(1 - \rho)$, and the total-compensation markdown wider by roughly 4 percentage points, evaluating (A.15) against (14) at median-decile estimates, about one point of which is the revaluation of amenity spending. Berger et al. (2026) report a similar widening, of about 3 percentage points, when top marginal rates rise in an oligopsony model with worker heterogeneity. The mechanism is similar: under a progressive schedule an additional unit of wages is taxed at a higher rate than the wage as a whole, so wage increases become a less effective way to attract

workers, and firms respond by paying less relative to marginal product.² Compensation in their model consists of wages only; in ours firms can also shift pay toward untaxed amenities, which reduces taxable wages and slightly offsets the elasticity decline, though both effects are small at our estimate of ρ .

Workers therefore bear more than the statutory burden of the tax. With the markdown wider, total compensation falls from roughly 0.79 to 0.75 of marginal product at the median decile: beyond the taxes they pay, workers receive a smaller share of output. Compensation also shifts toward untaxed amenities, which cushions workers' value but narrows the base. Both margins carry welfare costs at the firm level: employment, already below the value-taking benchmark, falls further, and the wedge $1/(1 - \tau)$ tilts the package toward amenities at any ρ , the standard distortion from exempting them from taxable income. In Berger et al. (2026) progressivity likewise widens markdowns, and the resulting misallocation of workers across firms is what lowers optimal progressivity there.

A.4 Model Extension: Production-Relevant Amenities

Some amenities, such as schedule flexibility or the possibility to work from home, affect not only worker utility but also output. To consider such amenities, we partition the amenity vector into *neutral* amenities, entering only utility, and *production-relevant* amenities $\{a_{jk}\}_{k \in \mathcal{P}}$, which also enter output through a transformation $p_k(a_{jk})$ with $p_k > 0$, $p'_k \leq 0$ and $p''_k \leq 0$: $Y = \psi_j \prod_{k \in \mathcal{P}} p_k(a_{jk}) L_j^\alpha$. Every amenity keeps its monetary cost $c_{jk} a_{jk} h_k(L_j)$, and neutral amenities are the case $p_k \equiv 1$. Providing one more unit of a production-relevant amenity now costs the firm its monetary price plus the output it forgoes, $(-p'_k/p_k)Y$, a real resource cost rather than a transfer. We write the sum as the effective marginal cost, $\Gamma_k \equiv c_{jk} h_k(L_j) + (-p'_k/p_k)Y$, and let $\mu_k \equiv c_{jk} h_k(L_j)/\Gamma_k$, with $\mu_k \in (0, 1]$ the cost share paid in money and $1 - \mu_k$ the share paid in output. An output-enhancing amenity such as training has $p'_k > 0$, hence $\mu_k > 1$; if all production-relevant amenities are of this type, $\theta < 1$ below and the comparisons with the canonical benchmark reverse. We present the output-reducing case. Because Γ_k contains output, the shares ω_k and σ_k , and with them $\mathcal{A}$, Φ and $\mathcal{M}$, now depend on ψ_j and the amenity levels as well as on L_j . For brevity, we suppress their arguments.

²In Berger et al. (2026) the elasticity a firm perceives adds to the extensive margin of which firm to work at (McFadden, 1974) an intensive margin of hours (MaCurdy, 1981), which they estimate to be small (a Frisch elasticity of 0.09, Marshallian elasticities near zero). Our model has only the extensive margin; adding an hours margin of up to 0.1 to the total-compensation elasticity in (14), which at the median decile is close to the wage elasticity of 3.6, would raise the implied total-compensation markdown by roughly half a percentage point, with hours and heads as substitutes in production.

Firm's problem. The firm's problem of Appendix A.1.2 changes in one place. The amenity FOC (A.2) carries Γ_k in place of $c_{jk}h_k(L_j)$, while the wage FOC (A.1) is unchanged, and dividing the two as in Appendix A.1.2 gives the MRS condition

$$\gamma_k \left(\frac{a_{jk}}{w_j} \right)^{\rho-1} = \frac{\Gamma_k}{L_j}. \quad (\text{A.16})$$

Multiplying (A.16) by a_{jk}/w_j and using $a_{jk}\Gamma_k = C_{jk}/\mu_k$ gives $\gamma_k a_{jk}^\rho = (\omega_k/\mu_k) w_j^\rho$ in place of $\omega_k w_j^\rho$ in (8): an amenity paid partly in output contributes more value than its monetary share, by the factor $1/\mu_k$. Summing across amenities, $V_j = (\mathcal{A}^\rho)^{1/\rho} w_j$ with the augmented multiplier $\mathcal{A}^\rho \equiv 1 + \sum_k \omega_k/\mu_k \geq \mathcal{A}$. $\mathcal{A}^\rho$ is what workers receive per dollar of wages and $\mathcal{A} = 1 + \sum_k \omega_k$ what the firm spends. Below, $\mathcal{A}$ always denotes the monetary multiplier, so $\mathcal{AM}$ captures monetary compensation against the marginal product, quantifiable from the spending data.

The wage FOC (A.1) balances the marginal revenue product against the wage, the marginal amenity bill, and the cost of raising job value for every inframarginal worker. Only the last term changes: labor supply responds to $V_j = (\mathcal{A}^\rho)^{1/\rho} w_j$, so the supply term $\mathcal{A}/\eta$ in the wage equation (9) becomes $\mathcal{A}^\rho/\eta$, while the marginal amenity bill $\sum_k \omega_k \varepsilon_L^{h_k}$ stays monetary. The wage markdown is therefore $\mathcal{M} = \eta/(\eta + \mathcal{A}^\rho + \eta \sum_k \omega_k \varepsilon_L^{h_k})$. To compare with the baseline (14), multiply by $\mathcal{A}$ and use the share identity $\mathcal{A}(1 - \Phi) = 1 + \sum_k \omega_k \varepsilon_L^{h_k}$. The “+1” in (14) then becomes $\theta \equiv \mathcal{A}^\rho/\mathcal{A} = 1 + \sum_k \sigma_k(1 - \mu_k)/\mu_k \geq 1$, so θ is the value delivered per unit of compensation that is paid out. Two results follow. The first concerns the markdown, which in the baseline equals the canonical benchmark $\eta/(\eta+1)$ at per-worker costs.

Proposition 1'' (Markdowns with Production-Relevant Amenities). *(i) The monetary-compensation markdown satisfies*

$$\mathcal{AM} = \frac{\eta}{\eta + \theta - \eta \Phi}, \quad (\text{A.17})$$

and crosses the canonical benchmark at a strictly positive curvature: $\mathcal{AM} \geq \eta/(\eta+1)$ iff $\Phi \geq (\theta-1)/\eta$. (ii) The wage markdown lies below it by the factor $\mathcal{A} > 1$, and the further gap $\mathcal{A}^\rho - \mathcal{A}$ is borne as forgone output, valued at the margin, rather than as payment. (iii) The canonical formula $\varepsilon_w^L/(\varepsilon_w^L + 1)$ no longer recovers the markdown in either case of Proposition 1(iii). All parts reduce to Proposition 1 when $\theta = 1$.

Forgone output thus acts like extra cost convexity of size $(\theta-1)/\eta$: the benchmark is reached only once concave costs offset the drag, at $\Phi = (\theta-1)/\eta$ rather than at zero. In

practice, a single production-relevant amenity with value share $\sigma_k = 0.02$ and monetary share $\mu_k = 0.2$, so that forgone output is four times the monetary bill, gives $\theta = 1.08$ and, at $\hat{\eta} = 4.3$, a crossing at $\Phi \approx 0.019$, just below the eighth decile's $\hat{\Phi} \approx 0.027$ in Section 5. At the baseline estimates, the set of deciles above the canonical benchmark is unchanged. A markdown below the benchmark therefore no longer signals convex costs. Nor can it be read off a supply elasticity (part (iii)): workers respond to $\mathcal{A}^P$ while the firm pays $\mathcal{A}$, and a supply elasticity measures the worker side only, so it cannot recover θ , which requires the output shares $1 - \mu_k$.

The second result revisits Proposition 2(i): the threshold cost elasticity $\varepsilon_L^{h_k}$ below which an amenity rises after a productivity shock. We derive it as the cost elasticity at which the amenity is exactly unchanged, $d \ln a_{jk} = 0$, where the output loss stays fixed. Parts (ii) and (iii) of Proposition 2 have no clean counterpart, since away from this point the output loss depends on the curvature of p_k . Log-differentiating the MRS condition (A.16) at $d \ln a_{jk} = 0$ gives $(1 - \rho) d \ln w_j = d \ln \Gamma_k - d \ln L_j$. At fixed a_{jk} , $d \ln \Gamma_k = \mu_k \varepsilon_L^{h_k} d \ln L_j + (1 - \mu_k) d \ln Y$ with $d \ln Y = d \ln \psi_j + \alpha d \ln L_j$. Substituting $d \ln L_j = \varepsilon_w^L d \ln w_j$ and $d \ln \psi_j$ from the log-differentiated wage equation (9) gives $\varepsilon_L^{h_k} - 1 = (\mu_k - \rho)/(\mu_k \varepsilon_w^L) + \frac{1 - \mu_k}{\mu_k} \varepsilon_L^M$, which at $\varepsilon_L^M = 0$ is the threshold in the following proposition.

Proposition 2'' (Productivity Shocks with Production-Relevant Amenities). *Let amenity k be the sole production-relevant amenity, let ε_w^L denote the elasticity of employment to the wage along the shock's path, and hold the markdown fixed ($\varepsilon_L^M = 0$). Following a positive productivity shock $d\psi_j > 0$, to first order the amenity level a_{jk}^* rises if and only if*

$$\varepsilon_L^{h_k} - 1 < \frac{\mu_k - \rho}{\mu_k \varepsilon_w^L}, \quad (\text{A.18})$$

that is, Proposition 2(i) with ρ replaced by ρ/μ_k . The two coincide at $\mu_k = 1$ and at $\rho = 0$.

The threshold falls under substitutes because the output part of the cost scales one-for-one with Y , while the worker's willingness to pay for the amenity scales as $w_j^{1-\rho}$, so cost outpaces benefit and the threshold drops below its neutral value $(1 - \rho)/\varepsilon_w^L$. Once $\rho > \mu_k$ it drops below the per-worker point, so even per-worker or concave-cost amenities can decline, which the baseline rules out. Under complements the shift reverses.

Illustration: flexible hours. The presence of production-relevant amenities can help interpret the decline in schedule flexibility found in Section 4.3 following a positive demand shock, without appealing to convex costs. Schedule flexibility plausibly has a small monetary share: the firm pays for it mostly in forgone output, through coverage and coordination,

rather than in money, so μ_k would be small. If $\mu_k < \rho$, which at $\hat{\rho} \approx 0.06$ means a monetary share below six percent, the right side of (A.18) is negative, so to first order the amenity rises only if its cost is concave enough, and per-worker flexibility would decline.³ In the limit $\mu_k \rightarrow 0$ it would decline for any cost curvature. This is consistent with what we report in Section 4.3: advertised flexible hours decline after a positive shock (Figure 6).

The same feature matters for measurement. An amenity paid in output widens the wedge between compensation and marginal product, $\mathcal{A}^P > \mathcal{A}$, without appearing in spending data. In the extreme where all amenities are of this type, $\mathcal{AM} = \mathcal{M} = \eta/(\eta + \mathcal{A}^P)$ while measured amenity spending is zero. Worker-side valuations can close the gap: estimates of willingness-to-pay price the value delivered however it is financed, so together with the expenditure-based multiplier they identify $\mathcal{A}^P - \mathcal{A}$ and hence the output shares. We leave this to future work.

A.5 Models of Imperfect Competition and Endogenous Amenities

Appendix Table A.2 summarizes models of imperfect labor market competition with endogenous amenity provision: the structure of amenity provision costs, the specification of preferences and the amenity bundle, the source of employer market power and whether the markdown each study derives depends on amenities, and the mapping into the amenity cost structure of our model. We depart from this literature by allowing multi-dimensional amenities with heterogeneous cost structures, and by focusing on the implications for the measurement of employer market power.

³The proposition holds the markdown fixed along the shock's path. If it moves, the threshold shifts by the second term in the expression derived above. In our estimates wages rise by less than value added per worker, so the markdown falls as the firm expands, $\varepsilon_L^M = 1 + 1/\varepsilon_w^L - 1/(\varepsilon_w^L \varepsilon_{VA}^w) \approx -0.3$ (Table 1), which lowers the threshold and reinforces this decline.

Appendix Table A.2: Overview of Models of Imperfect Labor Market Competition with Endogenous Amenity Provision.

Paper	Amenity	Worker Preferences	Provision Cost	Mapping to Our Model	Variation Used	Source of Market Power	Do Markdowns Depend on Amenities?
This paper	n endogenous amenities $\mathbf{a}$, plus exogenous $\bar{a}$	CES in $(w, \mathbf{a})$, $\hat{\rho} \approx 0.06$	$C_k = c_k h_k(L) a_k$, any size elasticity $\varepsilon_L^{h_k} \geq 0$	—	Firm-level demand shocks (Norway)	Idiosyncratic tastes	Yes: the wage and total compensation markdowns depend on amenity cost structure
Hwang et al. (1998)	Generic amenity a	Additive, $w + h(a)$	Per worker, C convex in a	One amenity, $\varepsilon_L^{h_k} = 1$; $\rho \rightarrow 1$ limit, but C is convex in a	—	Search frictions, wage posting	Not derived
Lang and Majumdar (2004)	Job characteristic h , e.g., hours	General $u(w, h)$	Forgone output	One production-relevant amenity (Appendix A.4); preferences not CES	—	Search frictions, wage posting	Not derived: profit per hire depends only on the utility offered
Flabbi and Moro (2012)	Hours flexibility, binary	Linear in wage and flexibility, worker-specific valuation	Share of match productivity	One binary amenity, $\varepsilon_L^{h_k} = 1$; $\rho \rightarrow 1$ limit; match-level bargaining	Structural estimation (US)	Search frictions, bargaining	Not derived: wages are Nash bargained with a fixed weight
Dube et al. (2022)	Vector of amenities (11 in the survey)	CES in the wage and a linear amenity index	Per worker	One composite amenity, $\varepsilon_L^{h_k} = 1$	Hypothetical job offers; Walmart’s minimum wage (US)	Idiosyncratic tastes	Yes: the Lerner rule for the wage has an amenity term
Arold et al. (2025), Appendix B	Worker rights in the contract, value $A(T)$	CES in after-tax wage and $A(T)$	Fixed, independent of firm size	One composite amenity, $\varepsilon_L^{h_k} = 0$; same CES; firm size exogenous	Provincial income tax rates (Canada)	Firm’s offer under a strike threat	Not derived
Bagga et al. (2025)	Remote work, binary	Linear, $w + x_i a$, two worker types	One-time random upgrade cost per job	One binary amenity; $\rho \rightarrow 1$ limit; no firm size	Aggregate amenity valuation and productivity shocks (US)	Search frictions, sequential auctions	Not derived
Colonna et al. (2025)	Working conditions, one amenity a	Cobb–Douglas in (w, a)	Per worker; mandated floor with fines	One amenity, $\varepsilon_L^{h_k} = 1$; $\rho = 0$ case	Random workplace inspections (Mexico)	Idiosyncratic tastes	Not derived: the amenity cost enters the wage first-order condition
Corradini et al. (2025), Appendix D	Female-friendly contract clauses	Stone–Geary in (w, a)	Per worker	One amenity, $\varepsilon_L^{h_k} = 1$; $\rho = 0$ case at zero reference levels	Union priorities (Brazil)	Idiosyncratic tastes	Yes: the markdown includes the amenity elasticity, as in Lagos (forthcoming)
Kroft et al. (2025), Supplemental Appendix H	Firm-specific amenity; allowed to respond to the shock	$\ln w + \ln g$, with g the value of amenities	Not modeled	One composite amenity, $\rho = 0$; their assumed gap is our $\varepsilon_w^L / \eta - 1$	Winning procurement auctions (US)	Idiosyncratic tastes	Yes: the wage-based markdown is biased if compensation and wages respond differently
Lavetti and Schmutte (2025)	Fatality risk R	$\ln w - h(R)$	Per worker, proportional to the wage	One amenity, $\varepsilon_L^{h_k} = 1$, $\rho = 0$; cost proportional to the wage	Job changers across industry–occupation cells (Brazil)	Idiosyncratic tastes, offer frictions	Not derived

continued on next page


Appendix Table A.2, continued.

Paper	Amenity	Worker Preferences	Provision Cost	Mapping to Our Model	Variation Used	Source of Market Power	Do Markdowns Depend on Amenities?
Sulka (2025)	Benefits b , binary, unobserved before hiring	Linear, $w + b$	Per worker, net of a tax advantage	One binary amenity, $\varepsilon_L^{h_k} = 1$; $\rho \rightarrow 1$ limit	—	Search frictions, wage posting	Not derived: workers receive their outside option
Adams-Prassl et al. (forthcoming)	Schedule flexibility f , amenity or disamenity	General $u(w, f)$	Profit term per worker, a gain or a cost	One production-relevant amenity (Appendix A.4), $\varepsilon_L^{h_k} = 1$; general preferences	Employer concentration (UK)	Cournot oligopsony, assumed supply curve	Yes: the wage is a markdown on output plus the amenity's profit term
Bíró et al. (2026)	Generic amenity a ; injury risk in the data	Linear, $w + a$	Per worker or fixed, convex in a	One amenity; $\rho \rightarrow 1$ limit; $\varepsilon_L^{h_k} = 1$ or 0, but C is convex in a	Local labor market tightness (Hungary)	Search frictions, wage posting	Not derived
Di Porto et al. (2026)	Low managerial delegation d	$v(w, d)$, $v_d \leq 0$	Forgone output	One production-relevant amenity (Appendix A.4); general preferences	Family vs. non-family firms (Italy)	Idiosyncratic tastes	Yes: the markdown depends on delegation through the supply elasticity
Faberman et al. (2026)	Generic amenity a	Linear, $w + \nu_g a$	Per job, C convex in a	One amenity, $\varepsilon_L^{h_k} = 1$; $\rho \rightarrow 1$ limit, but C is convex in a	Gender differences in amenity valuation (US)	Search frictions, sequential auctions	Not derived: wages are set by sequential auctions with a fixed bargaining weight
Lachowska et al. (2026)	Work hours h	$wh - \varepsilon_i h^\mu$	Forgone output	One production-relevant amenity (Appendix A.4); not nested	Hours differences across employers (US)	Exogenous firm surplus	No: the markdown is an exogenous firm-specific profit margin
Lagos (forthcoming)	Value index of all contract clauses	Stone–Geary in (w, a)	Per worker	One index, $\varepsilon_L^{h_k} = 1$; $\rho = 0$ case at zero reference levels	Union bargaining power (Brazil)	Idiosyncratic tastes	Yes: the markdown on total compensation includes the amenity elasticity
Lochner and Merkl (2026)	Flexibility requirement e	Linear, $w - e - \varepsilon_i$	Forgone output	One production-relevant amenity (Appendix A.4); $\rho \rightarrow 1$ limit	Job requirements over the business cycle (Germany)	Search frictions, bargaining	Not derived: wages are Nash bargained collectively with a fixed weight
Morchio and Moser (2026)	Generic amenity a	Linear, $w + \beta_g a$	Per worker, C convex in a	One amenity, $\varepsilon_L^{h_k} = 1$; $\rho \rightarrow 1$ limit, but C is convex in a	Gender differences in pay and worker flows (Brazil)	Search frictions, wage posting	Not derived: the amenity's net value is folded into firm productivity
Scott (2026)	Schedule variability and work outside regular hours	$\alpha \ln w + \beta' a$, by gender	Forgone output	Two production-relevant amenities (Appendix A.4); $\rho = 0$ case	Instruments from the firm's other jobs (France)	Idiosyncratic tastes	No: the markdown depends only on the wage elasticity

Notes: In our notation, a per-worker provision cost for amenity k corresponds to a size elasticity $\varepsilon_L^{h_k} = 1$, which implies $\Phi(L_j) = 0$ when this cost structure holds for every k . For brevity, this overview omits some models of search frictions with endogenous employer-provided health insurance. For instance, in Dey and Flinn (2005), insurance is bargained jointly with the wage at a premium per insured worker, while in Dizioli and Pinheiro (2016), Aizawa (2019), Fang and Shephard (2019), and Aizawa and Fang (2020), firms post a wage and an insurance offer and bear a fixed cost per firm, alone or in addition to the medical costs of their workers. None of these papers provide a markdown formula. We also omit models without employer market power in which amenities respond to demand. For instance, in Charles et al. (2022), price-taking firms without wage-setting power choose safety (via training that costs output and rental price that falls with output demand) against a constant per-accident cost. The last column refers to the markdown each paper states; *Not derived* means that the paper states none.


A.6 Numerical Illustrations for the Baseline Model

Appendix Figure A.2: Comparative Statics in Productivity ψ_j at $\rho = 0$ (Cobb–Douglas).




Notes: Within each panel, all lines are solid, and cost regimes are distinguished by color: concave ($\Phi(L_j) > 0$, blue), per-worker ($\Phi(L_j) = 0$, gray), convex ($\Phi(L_j) < 0$, red). Panels show (a) the wage w_j^* ; (b) the amenity multiplier $\mathcal{A}^*$; (c) the labor share $w_j^* L_j^* / R_j^*$ and (d) the total-compensation share $(w_j^* L_j^* + \sum_k C_{jk}^*) / R_j^*$ of revenue; and (e) the wage markdown $\mathcal{M}^* = w_j^* / \text{MRP}^*$ and (f) the total-compensation markdown $\mathcal{A}^* \mathcal{M}^* = \mathcal{A}^* w_j^* / \text{MRP}^*$. Single-amenity calibration with cost technology $h(L_j) = L_j^{\varepsilon_L^h}$ and parameters: $\eta = 5$, $\rho = 0$, $\alpha = 0.8$, $\gamma = 0.10$, $c = 1$, $N \bar{a}_j / \mathbb{V}^\eta = 20$, and $\varepsilon_L^h \in \{0, 1, 2.5\}$.

Appendix Figure A.3: Comparative Statics in Productivity ψ_j : Substitutes ($\rho > 0$).



Notes: Within each panel, all lines are solid, and cost regimes are distinguished by color: concave ($\Phi(L_j) > 0$, blue), per-worker ($\Phi(L_j) = 0$, gray), convex ($\Phi(L_j) < 0$, red). Panels show (a) the wage w_j^* ; (b) the amenity multiplier $\mathcal{A}^*$; (c) the labor share $w_j^* L_j^* / R_j^*$ and (d) the total-compensation share $(w_j^* L_j^* + \sum_k C_{jk}^*) / R_j^*$ of revenue; and (e) the wage markdown $\mathcal{M}^* = w_j^* / \text{MRP}^*$ and (f) the total-compensation markdown $\mathcal{A}^* \mathcal{M}^* = \mathcal{A}^* w_j^* / \text{MRP}^*$. Single-amenity calibration with cost technology $h(L_j) = L_j^{\varepsilon_L^h}$ and parameters: $\eta = 5$, $\rho = 0.3$, $\alpha = 0.8$, $\gamma = 0.10$, $c = 1$, $N\bar{a}_j / \mathbb{V}^\eta = 20$, and $\varepsilon_L^h \in \{0, 1, 2.5\}$.

Appendix Figure A.4: Comparative Statics in Productivity ψ_j : Complements ($\rho < 0$).



Notes: Within each panel, all lines are solid, and cost regimes are distinguished by color: concave ($\Phi(L_j) > 0$, blue), per-worker ($\Phi(L_j) = 0$, gray), convex ($\Phi(L_j) < 0$, red). Panels show (a) the wage w_j^* ; (b) the amenity multiplier $\mathcal{A}^*$; (c) the labor share $w_j^* L_j^* / R_j^*$ and (d) the total-compensation share $(w_j^* L_j^* + \sum_k C_{jk}^*) / R_j^*$ of revenue; and (e) the wage markdown $\mathcal{M}^* = w_j^* / \text{MRP}^*$ and (f) the total-compensation markdown $\mathcal{A}^* \mathcal{M}^* = \mathcal{A}^* w_j^* / \text{MRP}^*$. Single-amenity calibration with cost technology $h(L_j) = L_j^{\varepsilon_L^h}$ and parameters: $\eta = 5$, $\rho = -0.3$, $\alpha = 0.8$, $\gamma = 0.10$, $c = 1$, $N\bar{a}_j / \nabla^\eta = 20$, and $\varepsilon_L^h \in \{0, 1, 2.5\}$.

B Identification and Structural Estimation

This appendix describes the identification and estimation of structural parameters (η, ρ) .

B.1 Point Identification

Following the argument in Section 2.6, under non-Cobb–Douglas preferences ($\rho \neq 0$) and a non-degenerate cost-curvature index ($\Phi(L_j) \neq \bar{\Phi}$), both $\varepsilon_w^L(L_j)$ and $\varepsilon_L^A(L_j)$ vary across firms, and under a standard rank condition (at least two moments with sufficient variation in both), the cross-firm slope and intercept of (20) identify (η, ρ) , and then $\Phi(L_j)$ is identified from either elasticity. This argument requires the homogeneity of (η, ρ) across workers.

B.2 Partial Identification

The argument for point identification provided above breaks down if the cost-curvature index is degenerate or if preferences are Cobb–Douglas. We now characterize four configurations in which partial identification remains feasible despite the failure of point identification.

- (i) *Cobb–Douglas preferences* ($\rho = 0$) *with non-degenerate cost-curvature* ($\Phi(L_j) \neq \bar{\Phi}$). At $\rho = 0$, $\omega_k = \gamma_k$ regardless of c_{jk} , so $\bar{\mathcal{A}} = 1 + \sum_k \gamma_k$ is constant across firms. With $\Phi(L_j)$ varying, ε_w^L varies across firms while $\varepsilon_L^A = 0$, and the markdown is $\bar{\mathcal{A}}\mathcal{M}^* = \varepsilon_w^L/(\varepsilon_w^L + 1)$. Since $1/\varepsilon_w^L(L_j) = 1/\eta - \Phi(L_j)$ at $\rho = 0$, the profile of $\Phi(L_j)$ across firms is identified only up to the constant $1/\eta$. Under a sign restriction on cost elasticities, η is set identified. Appendix B.5 derives the set, which contains $\hat{\eta}$, and shows that it bounds the markdown decomposition at the bottom of the distribution.
- (ii) *Cobb–Douglas preferences* ($\rho = 0$) *with degenerate cost-curvature* ($\Phi(L_j) = \bar{\Phi}$). As in case (i), $\bar{\mathcal{A}}$ is constant across firms, and now both elasticities are too, with $\varepsilon_L^A = 0$ and the markdown $\bar{\mathcal{A}}\mathcal{M}^* = \varepsilon_w^L/(\varepsilon_w^L + 1)$. A constant multiplier $\bar{\mathcal{A}}$ across firms is consistent with $\rho = 0$ but does not establish it, since per-worker costs with common unit costs produce the same pattern at any ρ ; and under $\rho = 0$, η and $\bar{\Phi}$ cannot be separately identified. In the corner case $\bar{\Phi} = 0$, $\varepsilon_w^L = \eta$, as in the canonical monopsony.
- (iii) *Non-Cobb–Douglas preferences* ($\rho \neq 0$) *with zero cost-curvature* ($\Phi(L_j) = 0$). The multiplier $\bar{\mathcal{A}}_j$ varies across firms via c_{jk} but not within firms; both elasticities are constant, $\varepsilon_w^L = \eta$ and $\varepsilon_L^A = 0$, and the markdown is canonical, $\bar{\mathcal{A}}_j\mathcal{M}_j^* = \eta/(\eta + 1)$. Cross-firm variation in $\bar{\mathcal{A}}_j$, combined with $\varepsilon_L^A = 0$, pins down $\Phi(L_j) = 0$ and $\eta = \varepsilon_w^L$; ρ does not enter the elasticities and is not identified from them. If unit costs c_{jk} are observable across firms (e.g., from cost accounts), the cross-firm slope $d \ln \omega_k / d \ln c_{jk} =$

$-\rho/(1 - \rho)$, holding firm size fixed, identifies the sign and magnitude of ρ : a negative slope reveals substitutes ($\rho > 0$) and a positive slope complements ($\rho < 0$).

- (iv) *Non-Cobb–Douglas preferences* ($\rho \neq 0$) *with degenerate non-zero cost-curvature* ($\Phi(L_j) = \bar{\Phi} \neq 0$). The multiplier $\bar{\mathcal{A}}_j$ varies across firms via c_{jk} and within firms along the supply curve (since $\varepsilon_L^A \neq 0$). This is the only case with constant ε_w^L and $\varepsilon_L^A \neq 0$ (distinguishing it from case (ii)). The triplet $(\eta, \rho, \bar{\Phi})$ is not identified, but the markdown is pinned down as $\mathcal{A}^* \mathcal{M}^* = \varepsilon_w^L / [(\varepsilon_w^L + 1) + \varepsilon_w^L \varepsilon_L^A]$. The sign of ε_L^A further sign-identifies the primitives: $\text{sign}(\rho) = \text{sign}(\bar{\Phi})$ if $\varepsilon_L^A > 0$, and opposite signs if $\varepsilon_L^A < 0$; the supply-stability requirement $\varepsilon_w^L > 0$ adds the bound $\bar{\Phi} < (1 - \rho)/\eta$ when both are positive.

In each case the markdown $\mathcal{A}\mathcal{M}$ remains identified, whereas the data are consistent with $\rho = 0$ in cases (i) and (ii), pin down only η and $\Phi(L_j) = 0$ in case (iii), and none of $(\eta, \rho, \bar{\Phi})$ in case (iv). A non-zero ε_L^A at any firm rules out cases (i)–(iii), leaving case (iv) or the non-degenerate case, and in either its sign sign-identifies (ρ, Φ) as in case (iv).

B.3 Structural Estimation

For our structural analysis, we rely on evidence from the internal shock design. Specifically, we split the sample into deciles $d = 1, \dots, 10$ of baseline firm value added (residualized), and estimate the internal shock responses for each decile. The procedure yields, for every decile, post-shock average coefficients on log-employment, log-wage, and log-amenity multiplier, $(\hat{\beta}_L^d, \hat{\beta}_w^d, \hat{\beta}_A^d)$, together with their 3×3 covariance $\hat{\Omega}_d$. The event-study coefficients map into the decile-specific elasticities $\varepsilon_w^L(L_d) = \beta_L^d / \beta_w^d$ and $\varepsilon_L^A(L_d) = \beta_A^d / \beta_L^d$, since the demand-shifter scale drops out of the ratios. Substituting these into (20) and multiplying through by β_L^d gives the per-decile moment in linear β -form,

$$\beta_w^d = \frac{1}{\eta} \beta_L^d - \frac{1}{\rho} \beta_A^d.$$

We fit (η, ρ) by nonlinear least squares with parameter-dependent weights,

$$(\hat{\eta}, \hat{\rho}) = \arg \min_{(\eta, \rho)} \sum_{d=1}^{10} \frac{\hat{r}_d(\eta, \rho)^2}{\hat{v}_d(\eta, \rho)}, \quad (\text{B.1})$$

where $\hat{r}_d(\eta, \rho) = \hat{\beta}_w^d - (1/\eta) \hat{\beta}_L^d + (1/\rho) \hat{\beta}_A^d$ and $\hat{v}_d(\eta, \rho) = g' \hat{\Omega}_d g$ with $g = (-1/\eta, 1, 1/\rho)$. The variance $\hat{v}_d$ moves with (η, ρ) as $(\hat{\beta}_L^d, \hat{\beta}_A^d)$ are themselves estimated, and letting the weights move with the parameters delivers an errors-in-variables correction. Confidence intervals come from a parametric bootstrap: in each of $B = 2,000$ replications, we draw

$\beta^{*,d} \sim N(\hat{\beta}^d, \hat{\Omega}_d)$ per decile and the pre-shock multiplier $\mathcal{A}_d$ from its own standard error, re-fit (B.1), and store the parameter estimates and implied per-decile gradients of the cost-curvature index $\Phi(L_j)$, the total-compensation markdown $\mathcal{A}^*\mathcal{M}^*$, and the wage markdown $\mathcal{M}^*$. Reported 95% confidence intervals are the 2.5th and 97.5th percentiles of the bootstrap distribution.

B.4 Results: Preferences, Cost-Curvature and Markdowns

We first report the estimates of the structural preference parameters $(\hat{\eta}, \hat{\rho})$. With firms ranked by residualized baseline value added, we obtain $\hat{\rho} \approx 0.062$ with 95% bootstrap CI $[0.043, 0.085]$ and $\hat{\eta} \approx 4.3$ with 95% CI $[3.7, 5.3]$, implying $1/(1 - \hat{\rho}) \approx 1.07$ with CI $[1.05, 1.09]$. Ranking firms by residualized baseline employment instead gives $\hat{\rho} \approx 0.036$ with CI $[0.012, 0.071]$ and $\hat{\eta} \approx 4.6$ with CI $[3.4, 13.7]$. The over-identification test does not reject the constant- (η, ρ) restriction under either ranking: $Q \approx 12.6$ ($p \approx 0.13$) and $Q \approx 10.2$ ($p \approx 0.25$), each on 8 degrees of freedom. The parameter η is identified only through the intercept $1/\eta$ of the cross-decile restriction, where inverting a small, imprecise intercept yields the long right tail in the employment ranking's confidence interval.

Appendix Table B.1 reports the implied cost-curvature index $\Phi(L_j)$, total-compensation markdown $\mathcal{A}^*\mathcal{M}^*$, and wage markdown $\mathcal{M}^* = \mathcal{A}^*\mathcal{M}^*/\mathcal{A}_d$, with $\mathcal{A}_d$ the decile's pre-shock amenity multiplier, all evaluated at the corresponding $(\hat{\eta}, \hat{\rho})$: Panel A ranks firms by baseline value added, while Panel B by baseline employment. Under the value-added ranking, the implied cost-curvature moves from approximately -0.33 in the bottom decile to $+0.10$ in the top, with bootstrap confidence intervals that exclude zero in the bottom three deciles (negative) and in the top two (positive); the employment ranking gives a very similar profile, though its top-decile interval includes zero. The wage markdown rises from 0.59 to 0.77 (0.59 to 0.79 under the employment ranking). Since the pre-shock multiplier $\mathcal{A}_d$ varies far less across deciles than the elasticities (between 1.08 and 1.14 , with sampling error propagated in the bootstrap), it contributes only modestly to the wage-markdown gradient, which, like that of the total-compensation markdown, is driven by the cost-curvature wedge $\Phi(L_j)$.

B.5 Bounding the Cost-Curvature and Markdown Decomposition

The markdown decomposition in Section 5 recovers the level of $\Phi(L_j)$ in each decile from the multiplier elasticity $\varepsilon_L^A = \rho\Phi(L_j)/(1 - \rho)$, which is zero under Cobb–Douglas preferences, where case (i) above shows that $\Phi(L_j)$ is not point-identified. Requiring instead that no amenity have a total cost that falls with size, $\varepsilon_L^{h_k} \geq 0$, bounds $\Phi(L_j)$. The bounds reproduce the pattern of Section 5, with cost curvature dominating at the bottom of the distribution

Appendix Table B.1: Implied Distribution of Cost-Curvature and Markdowns.*Panel A: Ranked by Baseline Value Added.*

Decile	ε_w^L	$\Phi(L_j)$		$\mathcal{A}^*\mathcal{M}^*$		$\mathcal{M}^*$	
		Point	95% CI	Point	95% CI	Point	95% CI
1	1.721	-0.325	[-0.524, -0.171]	0.641	[0.572, 0.712]	0.585	[0.522, 0.649]
2	2.570	-0.146	[-0.246, -0.077]	0.725	[0.680, 0.760]	0.669	[0.627, 0.701]
3	3.079	-0.085	[-0.184, -0.011]	0.758	[0.711, 0.797]	0.701	[0.658, 0.737]
4	3.630	-0.039	[-0.126, 0.018]	0.786	[0.739, 0.822]	0.727	[0.683, 0.761]
5	3.641	-0.038	[-0.101, 0.011]	0.786	[0.754, 0.816]	0.726	[0.696, 0.753]
6	3.924	-0.020	[-0.078, 0.026]	0.798	[0.771, 0.822]	0.735	[0.710, 0.757]
7	4.055	-0.012	[-0.082, 0.042]	0.803	[0.770, 0.828]	0.737	[0.706, 0.759]
8	4.868	0.027	[-0.024, 0.070]	0.828	[0.802, 0.853]	0.755	[0.731, 0.778]
9	5.333	0.043	[0.001, 0.078]	0.840	[0.830, 0.849]	0.756	[0.747, 0.764]
10	7.548	0.095	[0.052, 0.130]	0.878	[0.860, 0.893]	0.772	[0.756, 0.785]

Panel B: Ranked by Baseline Employment.

Decile	ε_w^L	$\Phi(L_j)$		$\mathcal{A}^*\mathcal{M}^*$		$\mathcal{M}^*$	
		Point	95% CI	Point	95% CI	Point	95% CI
1	1.846	-0.314	[-0.582, -0.172]	0.654	[0.589, 0.714]	0.591	[0.533, 0.646]
2	2.447	-0.186	[-0.372, -0.073]	0.713	[0.660, 0.761]	0.652	[0.603, 0.696]
3	2.817	-0.134	[-0.324, -0.037]	0.741	[0.693, 0.780]	0.683	[0.639, 0.719]
4	3.233	-0.090	[-0.268, -0.011]	0.766	[0.717, 0.805]	0.706	[0.661, 0.742]
5	3.335	-0.081	[-0.257, -0.007]	0.771	[0.746, 0.795]	0.711	[0.687, 0.733]
6	4.039	-0.030	[-0.228, 0.057]	0.802	[0.754, 0.841]	0.739	[0.694, 0.775]
7	4.167	-0.023	[-0.194, 0.049]	0.807	[0.787, 0.827]	0.742	[0.723, 0.760]
8	4.649	0.001	[-0.164, 0.068]	0.823	[0.793, 0.848]	0.753	[0.725, 0.776]
9	5.445	0.031	[-0.136, 0.097]	0.844	[0.823, 0.862]	0.765	[0.746, 0.781]
10	8.143	0.090	[-0.070, 0.152]	0.888	[0.864, 0.907]	0.790	[0.770, 0.807]

Notes: $\varepsilon_w^L = \hat{\beta}_L^d / \hat{\beta}_w^d$ from the decile-specific event study; $\Phi(L_j) = (1 - \rho)(1/\eta - 1/\varepsilon_w^L)$, $\mathcal{A}^*\mathcal{M}^* = \eta/(\eta + 1 - \eta\Phi(L_j))$, and $\mathcal{M}^* = \mathcal{A}^*\mathcal{M}^*/\mathcal{A}_d$, with $\mathcal{A}_d$ the decile's pre-shock multiplier, at $(\hat{\eta}, \hat{\rho})$. Confidence intervals from the parametric bootstrap.

and $\hat{\eta}$ inside the admissible set, and a point decomposition at $\rho = 0$ under a restriction on how cost elasticities vary with size splits each gap almost as in our baseline.

Bounds on the cost-curvature index. At $\rho = 0$, the effective labor supply elasticity (12) is $\varepsilon_w^L(L) = \eta/(1 - \eta\Phi(L))$, so we have the following relationship for each decile d :

$$\frac{1}{\varepsilon_w^L(L_d)} = \frac{1}{\eta} - \Phi(L_d), \quad (\text{B.2})$$

where $\varepsilon_w^L(L_d)$ is the decile's estimated elasticity and $\Phi(L_d)$ the cost-curvature index at its observed average employment L_d in this decile. With $\varepsilon_L^{h_k} \geq 0$ for every amenity, $\Phi(L_d) = \sum_k \sigma_k(L_d)(1 - \varepsilon_L^{h_k})$ averages terms that are each at most one, so

$$\Phi(L_d) \leq s_d \equiv \sum_k \sigma_k(L_d) = 1 - \frac{1}{\mathcal{A}(L_d)},$$

the amenity share of total compensation, with equality when every amenity is a fixed cost. The lower bound is $\eta > 0$, that is $\Phi(L_d) > -1/\varepsilon_w^L(L_d)$. At that limit the firm has no wage-setting power, the cost-curvature factor $1/(1 - \Phi(L_d))$ equals the markdown, and the conduct wedge $1/(1 - \Phi(L_d)) - \mathcal{A}^*\mathcal{M}^*$ is zero. Solving (B.2) for $1/\eta$, we thus have

$$0 < \frac{1}{\eta} = \frac{1}{\varepsilon_w^L(L_d)} + \Phi(L_d) \leq \frac{1}{\varepsilon_w^L(L_d)} + s_d \quad \text{for every } d. \quad (\text{B.3})$$


The smallest of the ten caps binds. It is the top decile's, because $1/\varepsilon_w^L$ falls from 0.58 to 0.13 across deciles while s_d rises only from 0.088 to 0.120, which gives $\eta \geq 4$. The set is sharp, and each decile's cost curvature lies in $\Phi(L_d) \in (-1/\varepsilon_w^L(L_d), 0.253 - 1/\varepsilon_w^L(L_d)]$.

Point-identification at $\rho = 0$ under a parametric restriction. Cobb–Douglas preferences with common weights imply a constant amenity share, whereas the observed s_d rises from 0.088 to 0.120; we take the observed shares as given, since their variation is what identifies η below. For point-identification we add a restriction that our model does not impose: the share-weighted cost elasticity of the bundle depends on firm size in the same way for all firms, $1 - \bar{\varepsilon}(L) = f(L)$ with $f(L) = \phi_0 + \phi_1 \ln L$, so that $\Phi(L_d) = s_d f(L_d)$ and (B.2) becomes

$$\frac{1}{\varepsilon_w^L(L_d)} = \frac{1}{\eta} - s_d(\phi_0 + \phi_1 \ln L_d), \quad (\text{B.4})$$

which we estimate by least squares across the ten deciles; the intercept $1/\eta$ is identified because s_d varies across deciles. We obtain $\hat{\eta} = 4.1$, against 4.3 in the baseline, and a slope at which Φ rises by about 0.07 per log unit of employment.

Appendix Figure B.1: Bounds on the Decomposition of Total Compensation Markdown.




Notes: Three bars per decile. Bar I: decomposition under point identification at $(\hat{\eta}, \hat{\rho})$, as in Figure 8. Bar II: decomposition under point identification at $\rho = 0$ with the parametric restriction (B.4). Bar III: bounds at $\rho = 0$ under $\varepsilon_L^{h_k} \geq 0$; the orange area represents the cost curvature wedge for every admissible η , while the white area can be either a cost curvature wedge or a conduct wedge.

Results. Appendix Figure B.1 shows the decomposition under the baseline (bar I), under point identification at $\rho = 0$ with the parametric restriction (bar II), and under the bounds at $\rho = 0$ (bar III). Under any admissible η implied by the bound, cost curvature accounts for at least 67 percent of the gap between compensation and marginal product in the bottom decile, 43 percent in the second, 27 percent in the third, and a positive share through decile 6. The baseline puts these shares at 68, 46, and 32 percent. From decile 7 upward the bounds do not restrict the split, and the baseline attributes the gap there to conduct. The baseline $\hat{\eta} = 4.3$ lies inside the admissible set near its floor of 4.0. The implied decomposition using the parametric restriction (bar II) is close to our baseline at every decile.


C Appendix Figures and Tables

Appendix Figure C.1: The Effects of Demand Shocks on Total Value Added.




Notes: The main datasets used to construct outcomes are firm balance sheets. Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Eq. (21). Value-added outcomes are in logs. The internal shock (panel a) uses a balanced $[-3, +3]$ event window for cohorts 2005–2019 with standard errors clustered at the LLM and cohort level. The procurement shock (panel b) uses a balanced event window for cohorts 2012–2019 with standard errors clustered at the firm and cohort level. For these two shocks, we restrict to firms observed in all years of the balanced $[-3, +3]$ event window and having at least two stayers. The export shock (panel c) uses 2006 as the baseline period, and standard errors are clustered at the firm level. Stayers in this shock are defined as workers who are present at the firm for at least 8 years during the estimation window. The sample is restricted to firms with an export-to-revenue ratio of at least 1% in 2005–2007.

Appendix Figure C.2: The Effects of Internal Shock on Employment and Wage Margins.




(a) Employment Margins by Type




(b) Wage Margins by Type

Notes: The dataset used to construct outcomes is the matched employer–employee data (A-meldingen). Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Eq. (21). All outcomes are in logs. Wage outcomes are in hourly terms. We use the contracts sample (cohorts 2018–2019) for this internal shock and restrict this sample to firms observed in all years of the balanced $[-3, +3]$ event window, and with at least two stayers throughout the full estimation period. Estimates stack the 2018 and 2019 cohorts. Standard errors are clustered at the LLM level.


Appendix Figure C.3: Validation: Employment and Wages Impacts Across Samples.




(a) Employment: Internal Shock




(b) Wages: Internal Shock




(c) Employment: Procurement Sample



(d) Wages: Procurement Sample




(e) Employment: Export Sample



(f) Wages: Export Sample

Notes: The dataset used to construct outcomes is the matched employer–employee data. Coefficients are event-time estimates with 95% confidence intervals from the dynamic specification in Eq. (21). Panels (a) and (b) compare internal-shock estimates across the extended sample (cohorts 2005–2019) and the contracts sample (cohorts 2018–2019). Panels (c) and (d) re-estimate the internal shock restricted to the procurement-shock firms over the cohorts 2012–2019. Panels (e) and (f) re-estimate the internal shock using only the 2008 treatment cohort, restricted to export-shock firms in the 2002–2011 event window. Wages are defined for stayers. Standard errors are clustered at the firm level for the export shock, at the firm and cohort level for the procurement shock, and at the LLM and cohort level for the internal shock.

Appendix Figure C.4: Non-Wage Compensation and Non-Pecuniary Job Attributes.



Notes: Information on non-wage compensation (panels a–c) from firm cost accounts, 2015–2022. Information on contract types (panels d–f) from matched employer–employee data (A-meldingen), 2015–2022, and on advertised amenities (panels g–l) from job ads (NAV/FINN), 2015–2022, extracted using the text analysis methodology of Audoly et al. (2024). Bars report average rescaled residual values (panels a–c) and average shares among all employees (panels d–f) or all posted ads (panels g–l), respectively, within firm value-added quintiles (Q1–Q5), constructed as residuals from regressions on two-digit industry, commuting zone, and year fixed effects, rescaled to have a mean equal to the sample average of the raw outcomes. All monetary values are reported in logs and CPI-adjusted to 2015 Norwegian Kroner (NOK).

Appendix Table C.1: Variable Definitions and Construction.

Variable	Definition and Construction
<i>Firm Balance Sheets</i>	
Value Added	Log of value added from annual firm balance sheets (<i>Regnskapsregisteret</i>), constructed as gross revenues minus operating costs plus labor costs and depreciation, restricted to firms with positive values.
<i>Matched Employee–Employer Data</i>	
Employment	Log of the number of workers employed at the firm in a given year, drawn from annual matched employer–employee records, excluding workers with fewer than 52×4 annual contracted hours.
Mean Wages	Log of the average real hourly wage for all workers or stayers in the firm, constructed as annual earnings divided by contracted hours.
In-Kind Pay (<i>naturallytelser beløp</i>)	Log of the average monetary value of in-kind compensation for all workers or stayers in the firm, covering non-cash benefits such as free car access, company phone, lodging or meal coverage, employee loans, fitness-center access, and travel allowances, as reported by the employer to the tax authorities via <i>A-meldingen</i> .
In-Kind Contract (<i>andel ansatte som får naturalytelser</i>)	Firm-level share of workers or stayers receiving positive in-kind compensation from <i>A-meldingen</i> , defined as the count of stayers with in-kind pay above zero divided by the total number of workers or stayers.
Full-Time Contract (<i>andel ansatte med heltidsarbeid</i>)	Firm-level share of workers or stayers working more than 20 contracted hours per week, where total monthly hours sum contracted, variable, and overtime components from <i>A-meldingen</i> .
Fixed Hours Contract (<i>andel ansatte med fast arbeidstid</i>)	Firm-level share of workers or stayers without hours-based pay contracts (<i>timelønnsavtale</i>) or overtime hours (<i>overtidsarbeid</i>) recorded in <i>A-meldingen</i> .
Regular Daytime Contract (<i>andel ansatte med dagtidsarbeid</i>)	Firm-level share of workers or stayers on regular daytime work contracts (i.e., non-shift contracts), classified using the work-time arrangement variable from <i>A-meldingen</i> .
Permanent Contract (<i>andel ansatte med fast ansettelse</i>)	Firm-level share of workers or stayers on permanent contracts, classified using the employment contract variable from <i>A-meldingen</i> (available from 2021 onwards).

Continued on next page

Appendix Table C.1 – continued from previous page

Variable	Definition and Construction
<i>Firm Labor Cost Accounts</i>	
Wage Bill	Total wage payments to employees, recorded under entry P5000 of firms' annual business tax returns (<i>næringsoppgaven</i>), which include salaries, bonuses, and commissions. Used as the log of the total value and the per-worker value, and as the denominator in the amenity multiplier.
Pension Premiums (<i>pensjonskostnader</i>)	Total employer-sponsored pension premiums above the statutory minimum, recorded under entry P5420 of firms' annual business tax returns (<i>næringsoppgaven</i>). Used as a log of the total value, a log per worker, and a log as a share of the wage bill.
Other Non-Wage Personnel Costs (<i>andre personalkostnader</i>)	This measure includes non-wage personnel costs recorded under entry P5900 of firms' annual business tax returns (<i>næringsoppgaven</i>), covering staff gifts, subsidized canteen meals, courses, company cabins, housing, sports-club memberships, occupational and travel insurance, and mandatory pension contributions. The measure also includes employee travel and food reimbursements, coming from P7155, P7165, and P7370, which cover travel, per diem and mileage allowances, subsistence, and representation expenses. Used as a log of total costs, a log per worker, and a log as a share of the wage bill.
Amenity Cost Multiplier ($\mathcal{A}$)	Ratio of total labor costs to the wage bill from <i>Næringsoppgaven</i> , with $\mathcal{A} \equiv (\text{wage bill} + \text{total non-wage costs})/\text{wage bill}$, where total non-wage costs include pension premiums <i>P5420</i> and other non-wage personnel costs (<i>P5900</i> , <i>P7155</i> , <i>P7165</i> , and <i>P7370</i>). We winsorized at the 99th percentile; $\mathcal{A} \geq 1$ with $\mathcal{A} = 1$ by construction when non-wage costs are zero.
<i>Job Advertisements</i>	
On-the-Job Training	Computed as the share of the firm's postings in year t from NAV and FINN that mention apprenticeship, trainee programs, or on-the-job training, defined only for firms with at least one posting in year t .
Good Workplace	Computed as the share of the firm's postings in year t from NAV and FINN that mention at least one good-workplace attribute, including social environment, interesting tasks, good colleagues, work from home, open office, or employee involvement, defined only for firms with at least one posting in year t .

Continued on next page

Appendix Table C.1 – continued from previous page

Variable	Definition and Construction
Flexible Hours	Computed as the share of the firm’s postings in year t from NAV and FINN that mention convenient or flexible work hours, including flextime, standard hours, or no hour regulation, defined only for firms with at least one posting in year t .
Inconvenient Hours	Computed as the share of the firm’s postings in year t from NAV and FINN that mention inconvenient work hours, including shift work, non-standard hours, zero-hour contracts, and overtime, defined only for firms with at least one posting in year t .
Additional Perks	Computed as the share of the firm’s postings in year t from NAV and FINN that mention at least one additional perk, including proximity to nature, central location, exercise facilities, parking, company car, welfare benefits, company cabin, health benefits, subsidized food, or holiday benefits, defined only for firms with at least one posting in year t .
Good Finance	Computed as the share of the firm’s postings in year t from NAV and FINN that mention good financial benefits, including pension, insurance, or mortgage support, defined only for firms with at least one posting in year t .
Work from Home	Computed as the share of the firm’s postings in year t from NAV and FINN that mention the option to work from home, defined only for firms with at least one posting in year t .
Good Career Path	Computed as the share of the firm’s postings in year t from NAV and FINN that mention career development opportunities, defined only for firms with at least one posting in year t .

Notes: All monetary variables are deflated to 2015 NOK (CPI, base year 2015 = 100), and negative values are set to missing before taking logs. Matched employer–employee variables are firm-level outcomes from annual measures of *A-meldingen*, and contract-type variables are measured from a reference month, November. Firm labor cost account variables are from *Næringsoppgangen*, and each can enter our analysis in three forms: the log of the total firm-level value, the log per worker obtained by dividing by the number of workers reported in the matched employer–employee data, and the log as a share of the total wage bill. Job-ad amenity shares follow the text-analysis methodology of Audoly et al. (2024) applied to NAV and FINN postings, excluding staffing agencies, and are defined only for firms with at least one posting in year t .

Appendix Table C.2: Sample Sizes and Pre-Treatment Means Across Shocks.

Outcome	Internal Shock	Procurement Shock	Export Shock
<i>Firm Balance Sheets</i>			
Log Value Added	15.50	16.65	16.18
Number of Observations	3,817,030	375,883	36,759
Number of Firms	77,385	18,243	3,957
<i>Matched Employee–Employer Data</i>			
Log Employment	2.31	3.17	2.84
Log Wages	5.48	5.70	5.44
Number of Observations	3,817,030	375,883	36,759
Number of Firms	77,385	18,243	3,957
<i>Firm Cost Accounts</i>			
Log Amenity Cost Multiplier ($\ln \mathcal{A}$)	0.09	0.11	0.14
Number of Observations	3,714,409	372,770	33,463
Number of Firms	77,045	18,151	3,933

Notes: Pre-treatment means are computed over the pre-shock period ($t < c$). The table reports the number of firm-cohort-year observations and the number of unique firms used in each estimation. Wage outcomes are based on the stayers sample. Internal shock estimates use cohorts 2005–2019, procurement shock estimates use cohorts 2012–2019, and export shock estimates restrict the sample to firms with an export-to-revenue ratio of at least 1% in 2005–2007.

Appendix Table C.3: Industry Composition Across Estimation Samples.

Industry	Internal Shock	Procurement Shock	Export Shock
Agriculture	1.4%	1.0%	1.1%
Mining	0.5%	0.9%	2.3%
Manufacturing	9.6%	9.9%	41.5%
Electricity	0.4%	0.7%	0.1%
Water And Waste	0.4%	1.1%	0.9%
Construction	17.5%	26.1%	3.1%
Trade	28.0%	19.8%	37.6%
Transport	4.8%	5.2%	4.2%
Accommodation	5.0%	1.7%	0.0%
Information	4.5%	8.1%	2.5%
Finance	0.0%	0.0%	0.0%
Real Estate	3.1%	0.6%	0.4%
Professional	10.6%	12.8%	3.8%
Other Industries	14.0%	12.2%	2.7%
Number of Firms	77,385	18,243	3,957

Notes: This table reports the percentage share of treatment and control firms in the estimation sample belonging to that industry, classified according to the Norwegian Standard Industrial Classification (SN2007), which follows the European NACE Rev. 2 classification. The internal shock sample uses cohorts 2005–2019, procurement shock sample uses cohorts 2012–2019, and export shock sample restricts the sample to firms with an export-to-revenue ratio of at least 1% in 2005–2007. “Other Industries” bundle NACE sections N through S.

Appendix Table C.4: Internal Shock (Contracts Sample): Pre-Treatment Means.

Outcome	Pre-Treatment Mean	Obs.	Firms
<i>Firm Balance Sheets</i>			
Log Value Added	15.57	554,967	43,880
<i>Matched Employee–Employer Data</i>			
Log Employment	2.38	554,967	43,880
Log Wages	5.59	554,967	43,880
Log In-Kind Pay	6.39	387,194	33,749
Share In-Kind Contract	0.51	534,367	42,336
Share Full-Time Contract	0.91	534,367	42,336
Share Fixed Hours Contract	0.67	534,367	42,336
Share Regular Daytime Contract	0.99	534,367	42,336
<i>Firm Labor Cost Accounts – Totals</i>			
Log Wage Costs	15.16	551,109	43,634
Log Other Non-Wage Costs	12.14	543,821	43,409
Log Pension Premiums	11.62	302,876	28,976
Log Amenity Cost Multiplier ($\ln \mathcal{A}$)	0.09	551,159	43,638
<i>Firm Labor Cost Accounts – Per Worker</i>			
Log Wage Costs	12.78	551,109	43,634
Log Other Non-Wage Costs	9.76	543,821	43,409
Log Pension Premiums	9.10	302,876	28,976
<i>Job Advertisements</i>			
Share On-The-Job Training	0.24	150,461	21,114
Share Good Workplace	0.29	150,461	21,114
Share Flexible Hours	0.36	150,461	21,114
Share Inconvenient Hours	0.18	150,461	21,114
Share Additional Perks	0.20	150,461	21,114
Share Good Finance	0.13	150,461	21,114
Share Work from Home	0.00	150,461	21,114
Share Good Career Path	0.08	150,461	21,114

Notes: Pre-treatment means are computed over the pre-shock period ($t < c$). Obs. denotes the number of firm-cohort-year observations in the estimation, and Firms denotes the number of unique firms in the estimation. The outcomes with the matched employer–employee data, apart from employment, focus on the stayers sample. Numbers are for the stacked 2018 and 2019 cohorts.

Appendix Table C.5: Internal Shock (Contracts Sample): Post-Treatment Average Effects.

Outcome	$\bar{\beta}_Y$
<i>Firm Balance Sheets</i>	
Log Value Added	0.247*** (0.019)
<i>Matched Employee–Employer Data</i>	
Log Employment	0.160*** (0.016)
Log Wages	0.034*** (0.003)
Log In-Kind Pay	0.0195*** (0.0075)
Share In-Kind Contract	0.008*** (0.002)
Share Full-Time Contract	0.011*** (0.001)
Share Fixed Hours Contract	0.001 (0.002)
Share Regular Daytime Contract	0.0002 (0.0003)
<i>Firm Labor Cost Accounts – Totals</i>	
Log Wage Costs	0.193*** (0.015)
Log Other Non-Wage Costs	0.191*** (0.011)
Log Pension Premiums	0.173*** (0.023)
Log Amenity Cost Multiplier ($\ln \mathcal{A}$)	-0.00121*** (0.00038)
<i>Firm Labor Cost Accounts – Per Worker</i>	
Log Wage Costs	0.034*** (0.003)
Log Other Non-Wage Costs	0.031*** (0.009)
Log Pension Premiums	0.017* (0.010)
<i>Job Advertisements</i>	
Share On-The-Job Training	0.024*** (0.005)
Share Good Workplace	0.022*** (0.007)
Share Flexible Hours	-0.017*** (0.006)
Share Inconvenient Hours	0.007* (0.004)
Share Additional Perks	0.011 (0.010)
Share Good Finance	0.005 (0.004)
Share Work from Home	-0.0002 (0.0008)
Share Good Career Path	0.004 (0.005)

Notes: Reported coefficients correspond to the post-shock average effect $\bar{\beta}_Y = (\beta_{Y,1} + \beta_{Y,2} + \beta_{Y,3})/3$ from the dynamic specification in Eq. (21). Standard errors are reported in parentheses. The outcomes with the matched employer–employee data, apart from employment, focus on the stayers sample. The sample is restricted to firms observed in all years of the balanced $[-3, +3]$ event window and with at least two stayers throughout the full estimation period. We further restrict to firms with job postings for the outcomes with job ads. Estimates stack the 2018 and 2019 cohorts. Standard errors are clustered at the LLM level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table C.6: Post-Treatment Average Effects Across Shocks.

Outcome:	Internal Shock (1)	Procurement Shock (2)	Export Shock (3)
<i>Firm Balance Sheets</i>			
Log Value Added	0.260*** (0.016)	0.101*** (0.015)	0.293*** (0.052)
<i>Matched Employee–Employer Data</i>			
Log Employment	0.161*** (0.011)	0.079*** (0.011)	0.138*** (0.035)
Log Wages	0.044*** (0.003)	0.012*** (0.002)	0.074*** (0.016)
<i>Firm Cost Accounts</i>			
Log Amenity Cost Multiplier ($\ln \mathcal{A}$)	-0.00047* (0.00026)	-0.00150* (0.00078)	-0.00542 (0.00589)

Notes: Reported coefficients correspond to the post-shock average effect $\bar{\beta}_Y = (\beta_{Y,1} + \beta_{Y,2} + \beta_{Y,3})/3$, and for the export shock as $\bar{\beta}_Y = (\beta_{Y,2010} + \beta_{Y,2011})/2$, from the dynamic specification in Eq. (21). Standard errors are reported in parentheses. Wage outcomes are based on the stayers sample. Internal shock estimates use cohorts 2005–2019, procurement shock estimates use cohorts 2012–2019, and export shock estimates restrict the sample to firms with an export-to-revenue ratio of at least 1% in 2005–2007. Standard errors are clustered at the LLM and cohort level for internal shocks, at the firm and cohort level for procurement shocks, and at the firm level for export shocks. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix References

- Adams-Prassl, A., Balgová, M., Qian, M., and Waters, T. (forthcoming). “Firm Concentration and Job Design: The Case of Schedule-Flexible Work Arrangements”. *Review of Economics and Statistics*.
- Aizawa, N. (2019). “Labor Market Sorting and Health Insurance System Design”. *Quantitative Economics* 10 (4): 1401–1451.
- Aizawa, N. and Fang, H. (2020). “Equilibrium Labor Market Search and Health Insurance Reform”. *Journal of Political Economy* 128 (11): 4258–4336.
- Arold, B. W., Ash, E., MacLeod, W. B., and Naidu, S. (2025). *Worker Rights in Collective Bargaining*. NBER Working Paper 33605.
- Audoly, R., Bhuller, M., and Reiremo, T. A. (2024). *The Pay and Non-Pay Content of Job Ads*. Working Paper.
- Bagga, S., Mann, L., Şahin, A., and Violante, G. L. (2025). *Job Amenity Shocks and Labor Reallocation*. NBER Working Paper 33787.

- Berger, D. W., Herkenhoff, K. F., Jeong, J., and Mongey, S. (2026). *Labor Market Power with Worker and Firm Heterogeneity*. NBER Working Paper 35640.
- Bíró, A., Galindo Da Fonseca, J., Lindner, A., and Molnár, T. L. (2026). *Labor Market Tightness, Wage Inequality, and Workplace Amenities*. Working Paper.
- Charles, K. K., Johnson, M. S., Stephens Jr., M., and Lee, D. Q. (2022). “Demand Conditions and Worker Safety: Evidence from Price Shocks in Mining”. *Journal of Labor Economics* 40 (1): 47–94.
- Colonna, A., Pérez Pérez, J., and Rodrian, L. (2025). *Violation and Enforcement of Workplace Regulations: Evidence from Mexican Firm Inspections*. Working Paper.
- Corradini, V., Lagos, L., and Sharma, G. (2025). “Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs”. *Quarterly Journal of Economics* 140 (3): 2053–2105.
- Dey, M. S. and Flinn, C. J. (2005). “An Equilibrium Model of Health Insurance Provision and Wage Determination”. *Econometrica* 73 (2): 571–627.
- Di Porto, E., Pagano, M., Pezone, V., Saggio, R., and Schivardi, F. (2026). *Careers and Wages in Family Firms: Evidence from Matched Employer-Employee Data*. Working Paper.
- Dizioli, A. and Pinheiro, R. (2016). “Health Insurance as a Productive Factor”. *Labour Economics* 40: 1–24.
- Dube, A., Naidu, S., and Reich, A. D. (2022). *Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Walmart Workers*. NBER Working Paper 30441.
- Faberman, R. J., Mueller, A. I., and Şahin, A. (2026). *Job Search, Job Amenities, and the Gender Pay Gap*. IZA Discussion Paper 18418. IZA.
- Fang, H. and Shephard, A. J. (2019). *Household Labor Search, Spousal Insurance, and Health Care Reform*. NBER Working Paper 26350.
- Flabbi, L. and Moro, A. (2012). “The Effect of Job Flexibility on Female Labor Market Outcomes: Estimates from a Search and Bargaining Model”. *Journal of Econometrics* 168 (1): 81–95.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2017). “Optimal Tax Progressivity: An Analytical Framework”. *Quarterly Journal of Economics* 132 (4): 1693–1754.
- Holter, H. A., Krueger, D., and Stepanchuk, S. (2019). “How Do Tax Progressivity and Household Heterogeneity Affect Laffer Curves?” *Quantitative Economics* 10 (4): 1317–1356.
- Hwang, H., Mortensen, D. T., and Reed, W. R. (1998). “Hedonic Wages and Labor Market Search”. *Journal of Labor Economics* 16 (4): 815–847.
- Kline, P. (2025). “Labor Market Monopsony: Fundamentals and Frontiers”. *Handbook of Labor Economics* 6: 655–728.

- Kroft, K., Luo, Y., Mogstad, M., and Setzler, B. (2025). “Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry”. *American Economic Review* 115 (9): 2926–2969.
- Lachowska, M., Mas, A., Saggio, R., and Woodbury, S. A. (2026). “Work Hours Mismatch”. *Econometrica* 94 (3): 991–1025.
- Lagos, L. (forthcoming). “Union Bargaining Power and the Amenity-Wage Tradeoff”. *American Economic Review*.
- Lang, K. and Majumdar, S. (2004). “The Pricing of Job Characteristics When Markets Do Not Clear: Theory and Policy Implications”. *International Economic Review* 45 (4): 1111–1128.
- Lavetti, K. and Schmutte, I. M. (2025). *Estimating Compensating Wage Differentials with Endogenous Job Mobility*. Working Paper.
- Lochner, B. and Merkl, C. (2026). *Amenities and Wages over the Business Cycle: Dynamics and Gender Implications*. Discussion Papers in Economics 02/2026. Friedrich-Alexander-Universität Erlangen-Nürnberg.
- MaCurdy, T. E. (1981). “An Empirical Model of Labor Supply in a Life-Cycle Setting”. *Journal of Political Economy* 89 (6): 1059–1085.
- McFadden, D. (1974). “Conditional Logit Analysis of Qualitative Choice Behavior”. In: *Frontiers in Econometrics*. Ed. by Zarembka, P. Academic Press, pp. 105–142.
- Morchio, I. and Moser, C. (2026). “The Gender Pay Gap: Micro Sources and Macro Consequences”. *American Economic Review* 116 (5): 1765–1810.
- Qiu, X. and Russo, N. (2026). “Income Taxation across Countries”. *Economic Journal*.
- Scott, D. (2026). *Workplace Amenities and Labor Market Inequality*. Working Paper.
- Sulka, T. (2025). *Providing Benefits to Uninformed Workers*. CRC TRR 190 Discussion Paper 566. Rationality and Competition.