

Outsourcing, Inequality and Aggregate Output*

Adrien Bilal

Hugo Lhuillier

*Harvard University,
NBER & CEPR*

Princeton University

November 30, 2022

Abstract

Outsourced workers experience large wage declines, yet domestic outsourcing may raise aggregate productivity. To study this equity-efficiency trade-off, we contribute a framework in which multi-worker firms either hire imperfectly substitutable worker types in-house along a wage ladder, or rent labor services from contractors who hire in the same frictional labor markets. Three implications arise. First, selection into outsourcing: more productive firms are more likely to outsource to save on labor costs and higher wage premia. Second, a productivity effect: outsourcing leads firms to raise output and labor demand. Third, an outsourcing wage penalty: contractor firms pay lower wages. We find support for all three implications in French administrative data and rule out alternative explanations. Instrumenting revenue productivity using export demand shocks, we find evidence for selection into outsourcing. Instrumenting outsourcing using variation in occupational exposure, we find evidence for the productivity effect. We confirm the outsourcing wage penalty with a movers design. After structurally estimating the model and validating it against our reduced-form estimates, we find that the rise in outsourcing in France between 1997 and 2016 lowers low skill service worker earnings and welfare by 1.5%. Outsourcing increases labor market sorting, lowers the share of rents going to workers, but raises aggregate output by 6%. A simultaneous 5.5% minimum wage hike stabilizes earnings and maintains employment and output gains.

JEL Codes: E24, E25, E64, F12, J24, J31, J42, O40

*Adrien Bilal: adrienbilal@fas.harvard.edu. Hugo Lhuillier: hugo.lhuillier@princeton.edu. The authors thank Rodrigo Adao, Andres Blanco, Katarína Borovičková, Damien Capelle, Nik Engbom, John Grisgby, Emilien Gouin-Bonenfant, Gordon Hanson, Johnathan Hazell, Elhanan Helpman, Erik Hurst, Oleg Itskhoki, Gregor Jarosch, Larry Katz, Ilse Lindenlaub, Guido Menzio, Simon Mongey, Chris Moser, Tommaso Porzio, Stephen Redding, Richard Rogerson, Esteban Rossi-Hansberg, Rob Shimer, Stefanie Stancheva, Ludwig Straub, Nathan Zorzi, and numerous seminar participants for valuable discussions and comments. The authors gratefully acknowledge financial support from the Simpson Center for Macroeconomics, the International Economics Section, and the Industrial Relations Section at Princeton University. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the “Investissements d’Avenir” program (reference: ANR-10- EQPX-17—Centre d’accès sécurisé aux données—CASD). Special thanks to Xavier Ragot and the Observatoire Français des Conjonctures Économiques at Sciences Po Paris for offering physical access to the data.

Introduction

Domestic labor outsourcing is fundamentally changing the nature of the labor market. During the last two decades, firms have been increasingly concentrating on core competencies and contracting out a vast array of labor services, such as security guards, food and janitorial services. Workers in these occupations receive much lower wages at contractor firms than at traditional employers: the outsourcing wage penalty (Dube and Kaplan, 2010, Goldschmidt and Schmieder, 2017). This relative wage gap suggests that rising domestic outsourcing redistributes away from workers. At the same time, firms scale up more efficiently by contracting out. Outsourcing may thus generate aggregate productivity and employment gains that benefit workers. Despite the prevalence of outsourcing in the labor market, the tension between its distributional and productivity effects is far from understood. How does outsourcing shape aggregate production and its split between workers and firms?

The answer to this question depends on the fundamental driver of outsourcing. The *comparative advantage view* holds that contractor firms are more efficient at producing particular labor services because of gains from specialization in production. The comparative advantage view suggests that outsourcing raises aggregate Total Factor Productivity (TFP) as workers are reallocated towards more efficient contractors. Another perspective, the *cost-saving view*, holds that contractor firms simply enable their client firms to sidestep costly hiring and save on labor expenditures. The cost-saving view has more nuanced implications, as the reallocation of workers towards contractor firms may be neutral or even detrimental to aggregate TFP.

In this paper, we first build a theory of domestic outsourcing that disentangles the comparative advantage view from the cost-saving view. Second, we provide new reduced-form evidence of the distributional and productivity effects of outsourcing that our theory ties together using administrative data from France. Third, we structurally estimate our model and quantify the effects of outsourcing on inequality, rent-sharing and aggregate output.

Specifically, in the first part of the paper, we contribute a framework to study the emergence of outsourcing. We start with an environment that features three necessary ingredients, but no outsourcing yet. First, goods-producing firms are heterogeneous in productivity and have well-defined boundaries due to decreasing returns to scale in revenue. Second, not all workers are equally exposed to outsourcing. Firms hire workers of different skills who enter as imperfect substitutes in production.

Our third ingredient is key to rationalize the outsourcing wage penalty. Workers search for employment opportunities on and off the job along a wage ladder, and seemingly identical workers earn different wages at different employers. Crucially, as in Burdett and Mortensen (1998), scarce managerial time constrains hiring efforts. Managers have a limited span of control to recruit workers, and wages become an effective hiring tool. As a result, firms face an upward-sloping labor supply curve by skill as they internalize that higher pay attracts and retains more workers from competitors. Rent-sharing between workers and firms arises. We characterize wage and employment distributions in closed form. More productive firms with a larger target size pay higher wages. Wage inequality and dispersion in labor shares and markdowns across firms emerge in equilibrium.¹

We then introduce contractor firms in our environment. Contractor firms hire workers in the same

¹We use “wages” for concreteness but our theory applies equally to total compensation inclusive of benefits.

frictional labor markets and with the same recruiting technology as goods producers. Thus, they face the same upward-sloping labor supply curve. Instead of producing a consumption good, contractor firms sell labor services of their employees in a competitive labor service market. In the aggregate, contractor firms effectively expand managerial resources available for recruiting.

Consistently with the cost-saving view, goods producers may now outsource and buy labor services at an equilibrium price. In that case, goods producers bypass costly recruiting due to labor market frictions and constraints on managerial time. Instead, they indirectly tap into additional managerial resources by purchasing labor services in the competitive service market. When they outsource, goods producers switch from an upward-sloping labor supply curve to a vertical one. Decreasing returns to scale are thus critical to ensure that firm size remains well-defined in equilibrium.

Consistently with the comparative advantage view, contractor firms may be more productive at generating labor services than goods producers, for instance because of gains from specialization. Contractor firms may also be less productive than goods producers, for instance because of additional capital costs or communication and coordination costs from excess intermediation. The relative productivity of contractor firms captures the net balance between those forces.

Three main implications emerge. First, goods producers select into outsourcing. Productive firms who pay high wages to attract and retain a large workforce have the strongest incentives to outsource and save on labor costs. Less productive firms who pay lower wages prefer to hire in-house and avoid the additional cost of compensating contractor firms.

Second, outsourcing leads to a productivity effect at the firm level. When they outsource, goods producers scale up because their marginal cost of labor falls. Revenues and labor demand both rise. All else equal, outsourcing thus increases output and employment.

Third, outsourcing leads to a distributional effect. Its sign crucially hinges on whether the comparative advantage or cost-saving view holds. When contractor firms have a technological advantage, they pay higher wages than goods producers. By contrast, when contractor firms are technologically neutral or even disadvantaged, they pay lower wages than the marginal outsourcing firm.

Our theory is thus uniquely suited to discriminate between the comparative advantage and cost-saving views. When the law of one price fails as in our frictional economy, wages track the marginal product of labor. Thus, the comparative advantage view is incompatible with the outsourcing wage penalty that emerges in the data. By contrast, the cost-saving view in which goods producers outsource to save on labor costs is consistent with the outsourcing wage penalty. In that case, goods producers pass on the cost of labor to contractor firms that remain at a lower point on their labor supply curve.

We strengthen our results in two ways. First, we derive additional testable implications related to the comparative advantage and cost-saving views. Our micro-foundation of firm-specific labor supply curves implies that worker flows across firms are tied to the wage ordering between contractors and goods producers. Second, we show that our results continue to hold once we introduce firm-level idiosyncratic outsourcing costs to confront the predictions of our theory with the richness of the data.

To reach these conclusions, we have required a setup that departs from constant returns to scale and perfect substitution between workers. The wage-posting literature has imposed these assumptions ever since Mortensen and Vishwanath (1991) to retain traction. We overcome this technical challenge with two sufficient conditions. First, the revenue function exhibits a single-crossing property in firm

productivity and employment. This condition ensures that more productive firms always prefer to hire more and nests most standard revenue functions. Our second condition consists in a trembling-hand equilibrium refinement that precludes non-smooth equilibria.

In the second part of the paper, we test the novel implications of our theory using administrative data from France and rule out alternative mechanisms. We combine matched employer-employee data from employer tax returns, balance sheet records for the universe of firms, firm-level customs data and a firm-level survey that details outsourcing information. We measure outsourcing at the firm level as expenditures on external workers: workers who are not employees of the firm, but are at least partially under the legal authority of the purchasing firm. We identify contractor firms using industry codes. Our main analysis starts in 1996 and stops in 2007 due to a change in data collection procedures. Aggregate expenditures on outsourcing represent 6% of the aggregate wage bill in 1996 before rising to almost 11% in 2007. Extrapolating beyond 2007 suggests that outsourcing may represent between 11% and 20% of aggregate wages by 2016.

We first examine the distributional effect of outsourcing. Consistently with the cost-saving view, we show that contractor firms locate at the bottom of the job ladder along a number of labor market statistics. We start by confirming the outsourcing wage penalty first documented by Dube and Kaplan (2010) and Goldschmidt and Schmieder (2017), but in French data. Contractor firms pay wages that are 14% below other firms after controlling for individual worker heterogeneity. In line with our microfoundation for monopsony power, we show that contractors also hire less from employment than other firms, churn more through workers, and have negative net poaching.

Second, we demonstrate that firms select into outsourcing. We ask whether firms with large value added have a high outsourcing share as in our model. We define the outsourcing share as outsourcing expenditures out of all labor costs including outsourcing. We find a robust correlation between the outsourcing share and value added across and within firms. As highlighted by our extension with idiosyncratic outsourcing costs, a causal interpretation of this relation may be confounded by unobserved shocks such as improvements in Information Technologies (IT) at the firm level. To isolate the effect of a change in revenue productivity, we instrument for it with shift-share export demand shocks at the firm level: we interact initial exports shares of exporters across destination markets with changes in foreign demand (Hummels et al., 2014). Our identification assumption is that changes in foreign demand are unrelated to the idiosyncratic outsourcing cost of any particular firm (Borusyak et al., 2021). Our Two Stage Least Squares (2SLS) estimate implies that a 10% increase in value added driven by quasi-exogenous foreign demand shocks leads to a 0.33 percentage points (p.p.) rise in the outsourcing share.

Third, we provide evidence for the productivity effect. We show that a decline in the cost of outsourcing leads firms to scale up. Our model indicates that we need to isolate declines in firm-level outsourcing costs from confounding changes in revenue productivity. To that end, we instrument the outsourcing share with shift-share outsourcing shocks using variation in exposure at the firm level: we interact firm-level initial occupation shares with changes in average outsourcing expenditures at the occupation level. Our identification assumption is that initial occupation shares are unrelated to subsequent changes in revenue productivity (Goldsmith-Pinkham et al., 2020). Our 2SLS estimate indicates that a 1 p.p. rise in the outsourcing share driven by quasi-exogenous outsourcing cost shocks

leads to a 7.9% rise in value added.

We rule out several alternative mechanisms that could explain our findings. Outsourcing could arise because it provides more flexibility and helps firms insure against demand volatility. We show that more volatile industries use less outsourcing, not more.² Firms could outsource to alleviate equity concerns. We show that firms with more unequal pay structures outsource less, not more.³ Selection into outsourcing could emerge in part because of union wage-setting or size-dependent policies. With a regression discontinuity design, we show that the major threshold (50 employees) for union participation and policies in France indeed generates bunching, but no detectable jump in outsourcing.

In the third part of the paper, we develop and structurally estimate a quantitative version of the framework. The main addition is a flexible curvature in the recruiting cost function of goods producers and contractor firms. This extension captures that firms may expand their in-house human resources department at some cost and hire in-house without raising wages too rapidly. Our cost-saving channel of outsourcing then operates as long as there is some curvature in this cost function. We also let contractors have a comparative advantage in recruiting to match their relative size together with the outsourcing wage penalty. We use three skill types. High skill and core low skill workers are shielded from outsourcing. Only low skill service workers are exposed to outsourcing.

We estimate the model with a Method of Simulated Moments (MSM) estimator and target cross-sectional moments only. We infer from outsourcing expenditure and employment shares that workers at contractor firms are 35% less efficient than in-house workers. This result falsifies the comparative advantage view but is consistent with communication and monitoring frictions lowering the efficiency of outsourced relative to their in-house counterparts. We also infer that contractors are twice as efficient at recruiting workers than goods producers from their relative size. We interpret this result as the core activity of contractors being to screen and recruit workers. Thus, the data favors a view in which contractor specialize in recruiting activities rather than in production activities.

We support our inference by confirming that our estimated model matches non-targeted moments. It reproduces our reduced-form within-firm estimates for selection into outsourcing and the productivity effect. It also replicates residual wage inequality and differences in worker flows between contractors and goods producers. The distribution of firm-level labor shares is a key metric of rent-sharing and lines up in the model and in the data.

We then quantify the race between the productivity and the distributional effects of outsourcing in the aggregate. Our main counterfactual changes both the demand and supply of outsourcing such that outsourcing expenditures track the rise in France between 1997 and 2016.⁴ We estimate demand and supply shocks by matching industry-level panel regressions of value added, wages and employment onto the outsourcing share using data from 1997 to 2007. Supply shocks such as improvements in the recruiting technology of contractor firms rationalize simultaneous employment increases and wage declines. Demand shocks additionally rationalize rising value added. We interpret this combination of demand and supply shocks as improvements in information technology that facilitate standardization

²At the firm level, more volatility is positively correlated with outsourcing, but the relationship is three times weaker than the effect of firm scale that we highlight.

³Our results are consistent with equity concerns (e.g. Card et al., 2012, Breza et al., 2017) since equity concerns primarily bind for workers within the same occupation, rather than across occupations.

⁴We use 15% as our 2016 target as it lies in the middle of the range from our extrapolation.

and monitoring of tasks between goods producers and contractor firms.

Our first result is that low skill service workers are worse off because of outsourcing. We focus on changes in expected earnings as they coincide with welfare in our environment. We find that expected earnings of low skill service workers fall by nearly 1.5% by 2016. We unpack this result into partial and general equilibrium channels.

The first, partial equilibrium channel is that service workers are reallocated towards contractor firms, where their employment share rises by 21 p.p. by 2007. Despite being non-targeted, this large reallocation tracks the data reasonably well (16 p.p.). Together with the 14% outsourcing wage penalty, this reallocation drives a 1.5% decline in expected earnings for low skill service workers by 2007. By 2016, earnings fall by 2.8%.

General equilibrium channels partly offset this decline. The productivity effect materializes as substantial employment gains for low skill service workers, leading to a 5% rise in earnings by 2016. Further, as contractors expand, competition at the bottom of the job ladder increases and pushes up wages at low-paying firms. Yet, outsourcing rarefies high-paying job opportunities, and expected earnings decline by another 3.8%. Our decomposition highlights that reduced-form approaches that only pick up the first, partial equilibrium impact miss key general equilibrium adjustments and overstate welfare losses of service workers.

In the aggregate however, the economy benefits from outsourcing. Output rises by 6.1% due to the productivity effect. This increase is entirely driven by employment gains. Perhaps surprisingly, Total Factor Productivity (TFP) declines by 2%. We use an exact TFP decomposition to understand this reduction. In the aggregate, outsourcing increases TFP by 7% holding effective labor fixed: outsourcing increases allocative efficiency among goods producers by increasing labor demand at firms that were most constrained by labor market frictions and had a high marginal product of labor. Yet, consistently with the cost-saving view, outsourcing also lowers TFP by 9% by reallocating workers towards less efficient contractors. By contrast, the comparative advantage view would imply that both effects are positive, crucially overstating aggregate TFP gains.

Our environment is uniquely equipped to connect the welfare and productive effects of outsourcing to labor market sorting and inequality (Song et al., 2018). We find that outsourcing leads to rising labor market sorting but relatively stable wage inequality. The reallocation of low skill service workers towards low-paying contractors increases labor market sorting. The correlation between worker and firm wage premia rises by 23 p.p. in the model. Despite this increase, the pro-competitive effects of contractors raise goods producer wages for service workers at the bottom of the job ladder by over 10%. Thus, the wage gap between low skill service and high skill workers conditional on the same employer shrinks. This general equilibrium feedback offsets the impact of sorting on wage inequality which remains stable.

Inequality between workers and firm shareholders rises, however, as the share of rents going to workers shrinks. In our monopsony environment, the share of the marginal product of labor that workers receive—the markdown—is determined in equilibrium.⁵ Markdowns fall for service workers as the highest paying firms switch to outsourcing. In principle, core and high skill workers could still benefit from outsourcing because of complementarities in production and the increase in service

⁵We specify the close connection between markdowns and firm-level labor shares in Section 5.3.

labor demand. Nevertheless, despite an increase in their marginal product of labor, core and high skill workers do not gain from outsourcing because their markdowns fall too. As a result, firm profits rise by 2%. Our results highlight the key role of equilibrium adjustments in markdowns in the transmission of productive gains from outsourcing to workers.

We conclude our paper by asking whether simple labor market policies can mitigate expected earnings losses due to worker reallocation while maintaining most of the employment and output gains. A 5.5% increase in the minimum wage between 1997 and 2016 ensures that expected earnings of low skill service workers remain constant. Output gains are nearly identical to the baseline increase. Thus, minimum wage reforms can ensure that outsourcing benefits low skill service workers while preserving the productivity effects of outsourcing.

This paper relates to several strands of literature. The first is the rapidly expanding empirical literature that studies the distributional and productivity effects of outsourcing. Dube and Kaplan (2010), Goldschmidt and Schmieder (2017), Dorn et al. (2018) and Drenik et al. (2020) document that domestic outsourcing is on the rise and that outsourced workers experience wage declines in the U.S., Germany and Argentina, respectively. Abraham and Taylor (1996) provide an early discussion of the cost-saving and comparative advantage views.⁶ Segal and Sullivan (1997), Katz and Krueger (2017) and Katz and Krueger (2019) document a rise in alternative work arrangements in the U.S.⁷ Bertrand et al. (2020) show that an increase in the supply of contract labor helped Indian firms scale up, and Bostanci (2021) highlights the tension between trade secret protection and the productivity effect of outsourcing. Munoz (2022) documents the related role of posted workers, although their employment share of 0.5% remains well below that of domestic contractors of 5-9%. We contribute to this literature by providing a unified general equilibrium theory of outsourcing in which we disentangle the comparative advantage from the cost-saving views, tie the productivity and distributional effects together, and analyze the trade-off between both forces in the aggregate.⁸

Second, our paper relates to the outsourcing literature that focuses on the make-or-buy choice that firms face (Grossman and Hart, 1986, Hart and Moore, 1990, Grossman and Helpman, 2002). Our theory defines the boundary of the firm in product markets through decreasing returns, but requires that goods producers cannot take ownership of contractor firms as a whole. By abstracting from the particular frictions at work in ownership markets, our theory uniquely delivers specific interpretations of the labor market consequences of domestic outsourcing that have been recently documented, together with aggregate efficiency implications.

Third, our paper connects to the literature studying how labor market frictions give rise to factor price dispersion and misallocation. We contribute to the wage-posting monopsony tradition (Burdett

⁶We contrast our results for selection into outsourcing and demand volatility with the ones in Abraham and Taylor (1996) in Sections 3.4 and 3.6 respectively.

⁷Giannoni and Mertens (2019) emphasize the impact of outsourcing on the labor share in the U.S. Bergeaud et al. (2020) highlight that internet broadband expansion leads firms to concentrate on their core activities in France. Relatedly, LeMoigne (2020) highlights that the consequences of fragmentation events for workers resemble those of outsourcing events. See Handwerker (2021) for a similar idea in the U.S. and Bernhardt et al. (2016) for a review of the earlier literature on outsourcing in the U.S.

⁸Since this paper was first circulated, two papers subsequently emphasized related points. Spitze (2022) studies wages and benefits for low and high skill outsourced workers in the U.S. Felix and Wong (2021) provide evidence for the productivity effect in Brazil.

and Mortensen, 1998, Berg and Ridder, 1998, Bontemps et al., 2000, Engbom and Moser, 2021, Heise and Porzio, 2022) by providing sufficient conditions to depart from constant returns to scale and perfect substitutability between workers in production which has remained an open problem since at least Mortensen and Vishwanath (1991). As in Gouin-Bonenfant (2022), our monopsony environment delivers dispersion in markdowns and labor shares across firms that is consistent with evidence in Autor et al. (2020).⁹

Finally, our paper relates to the literature on trade in intermediate inputs and international offshoring (Feenstra and Hanson, 1999, Antràs, 2003, Grossman and Helpman, 2005, Grossman and Rossi-Hansberg, 2008, Acemoglu et al., 2015, Antràs et al., 2017). When firms trade intermediate inputs, they contract on a physical good. When firms outsource domestically, they contract on the flow of services of a worker, thereby leading to distinct implications for wage inequality. When firms offshore internationally, they take advantage of lower wages in other countries. Domestic outsourcing reflects similar forces, but requires first to break the law of one price in the domestic labor market.

The rest of this paper is organized as follows. Section 1 lays out the basic framework without outsourcing. Section 2 introduces outsourcing in the economy. Section 3 details the reduced-form results supporting our theory. Section 4 lays out the quantitative extensions of the model and the structural estimation. Section 5 presents our counterfactuals. The last section concludes. Proofs and further details can be found in the Appendix and the Online Appendix.

1 A theory of wage premia with large firms

1.1 Setup

Time is continuous, and we focus on a steady-state equilibrium. There is a unit measure of workers. Each worker is characterized by its exogenous and permanent skill type $s \geq 0$. Types are distributed in the population according to the measure $m_s ds$ with respect to a base measure denoted by ds .¹⁰ Workers have linear preferences in income, inelastically provide one unit of labor per time period, and discount future utility at rate r . They can be either employed or unemployed, in which case they earn skill-specific unemployment benefits b_s .

A measure $M^G > 0$ of goods-producing firms populates the economy. Firms are indexed by productivity z with support $[\underline{z}, \bar{z}]$. The corresponding cumulative distribution function Γ admits a finite and continuous density. For simplicity, \underline{z} is large enough relative to $\sup_s b_s$ so that all matches are viable. A firm with productivity z that hires a measure n_s of workers of each skill s generates revenue $R(z, \mathbf{n})$, where $\mathbf{n} = \{n_s\}_s$ denotes the vector of employment across worker types. R is twice continuously differentiable and increasing in each argument.

⁹A recent influential literature has developed models based on extreme-value distributed compensating differentials (Card et al., 2018, Sorkin, 2018, Berger et al., 2022, Lamadon et al., 2022). These frameworks are tractable and relevant to study market-level effects of inequality and mobility. This tractability comes with two caveats that are not well-suited to study outsourcing. These models imply welfare equalization across firms within markets, making the connection between inequality and welfare difficult. They also generate constant within-market markdowns and labor shares under standard functional forms, making it difficult to rationalize any dispersion in labor shares across firms. We show that this dispersion is key to understand the impact of outsourcing on inequality in Section 5.3. See Online Appendix F.2 for a detailed discussion.

¹⁰This notation allows us to capture both continuous and discrete type distributions without loss of generality.

Labor markets are segmented by skill s . Unemployed workers of skill s sample wage offers randomly at Poisson intensity λ_s^U . Employed workers of skill s sample wage offers with intensity $\lambda_s^E \leq \lambda_s^U$ from the same distribution. Employed workers can break their current contract to accept a new wage offer. Existing matches are destroyed at Poisson rate δ_s .

Firms optimally post wage offers in every skill-specific labor market to attract and retain workers. As in Burdett and Mortensen (1998), firms commit to a single, fixed, and non-state-contingent wage by skill. Wages cannot be renegotiated throughout employment spells. Every firm is endowed with a unit measure of managerial time to devote to recruiting activities (“vacancies”) for every skill s to which they attach the same skill-specific wage offer.

1.2 The labor supply curve

To understand the skill-specific labor supply curve faced by each firm, we must first characterize the job search behavior of workers. This subsection follows closely Burdett and Mortensen (1998). Given the equilibrium distribution of wage offers for skill s , denoted $F_s(w)$, the value of unemployment and the value of being employed at a given wage w satisfy:

$$\begin{aligned} rU_s &= b + \lambda_s^U \int \max\{V_s(w) - U_s, 0\} dF_s(w) \\ rV_s(w) &= w + \lambda_s^E \int \max\{V_s(w') - V_s(w), 0\} dF_s(w') + \delta_s(U_s - V_s(w)). \end{aligned}$$

The value of being employed at wage w , $V_s(w)$, is increasing with the wage w , so that workers behave as income maximizers: they always accept higher wage offers while employed.

The movement of workers up the job ladder determines the skill-specific labor supply curve faced by each firm. To characterize it, we solve for the equilibrium distribution of wages of employed workers $G_s(w)$. By equating inflows and outflows of workers in each wage interval, we relate the wage offer distribution $F_s(w)$ to the wage distribution of employed workers $G_s(w)$ in Appendix A.1. We obtain

$$G_s(w) = \frac{F_s(w)}{1 + k_s(1 - F_s(w))}, \quad k_s = \frac{\lambda_s^E}{\delta_s}. \quad (1)$$

From equation (1) we characterize the number $N_s(w)$ of employed workers per wage offer w for every skill s :

$$N_s(w) = \frac{(1 + k_s)e_s}{(1 + k_s(1 - F_s(w)))(1 + k_s(1 - F_s(w^-)))}, \quad (2)$$

where $e_s = \frac{\lambda_s^U m_s}{\delta_s + \lambda_s^U}$ is the measure of employed workers of skill s , and $F_s(w^-)$ denotes the left-limit of F_s at w .

Crucially, the labor supply curve $N_s(w)$ is non-decreasing in the wage w , with a slope that depends on the equilibrium distribution of wage offers in the economy, $F_s(w)$. We turn to the decision problem of firms to characterize this distribution.

1.3 Wage and employment distributions

The number $n_s(w)$ of workers per firm posting wage w is simply related to the number of workers employed at every wage by $n_s(w) = N_s(w)/M^G$ since firms have a unit measure of managerial time to devote to recruiting. When the discount rate is low enough, firms choose their wage offers $\{w_s(z)\}_s$ to maximize their flow profits.¹¹ Firms internalize their skill-specific labor supply curves: they attract and retain more workers if they pay more. Flow profits are given by:

$$\begin{aligned} \pi(z) &= \max_{\{w_s, n_s\}_s} R(z, \{n_s\}_s) - \int w_s n_s ds \\ \text{s.t. } n_s &\leq n_s(w_s) = \frac{N_s(w_s)}{M^G}. \end{aligned} \quad (3)$$

Unless the distribution $F_s(w)$ can be characterized more precisely, the problem in equation (3) is intractable in general equilibrium. The wage-posting literature—from Burdett and Mortensen (1998) to Engbom and Moser (2021)—has leveraged a key simplifying assumption to make progress. Under constant returns and perfect substitutability of workers in production, $R(z, \mathbf{n}) = z \int n_s ds$, the problem (3) can be split at the match level. Once decoupled across matches, it is straightforward to see that (3) exhibits a single-crossing property. This structure implies that wages are increasing in productivity z , which in turn allows to solve for the distribution of wage offers in terms of the equilibrium wage policy and the exogenous productivity distribution, $F_s(w_s(z)) = \Gamma(z)$.

Outsourcing however requires a well-defined boundary of the firm as well as possible interactions between workers in production. Handling these features has remained an open problem since Mortensen and Vishwanath (1991) who pointed out that the usual arguments for monotone wages and uniqueness do not apply. We overcome the challenges that come with this departure from linearity with two sufficient conditions. Our first and main sufficient condition imposes minimal structure on the revenue function R that lets us rank wages by firm productivity.

Assumption (A). $(z, \mathbf{n}) \mapsto R(z, \mathbf{n})$ is strictly supermodular in all its arguments.

Given that R is twice continuously differentiable, Assumption (A) is equivalent to imposing strictly positive cross-derivatives between all arguments. It amounts to a form of complementarity between productivity and every labor type, as well as between any two types of labor. Assumption (A) ensures that more productive firms prefer to hire more workers of every type.

Importantly, the complementarities built in Assumption (A) stand in productivity and employment levels, as opposed to the usual notion of complementarity between worker types that stands in proportions. Our supermodularity assumption is thus compatible with a wide class of revenue functions and allows for workers to be complements or substitutes in production in the usual sense.

For instance, consider the revenue function $R(z, \mathbf{n}) = z \left(\int (a_s n_s)^{1-\frac{1}{\eta}} ds \right)^{\frac{\eta(\sigma-1)}{\sigma(\eta-1)}}$. Such a revenue function arises when workers have CES demand over M^G differentiated varieties with elasticity of substitution $\sigma > 1$, and firms produce with a CES production function with elasticity of substitution η between skills. This revenue function also arises if there are technological decreasing returns to scale in production. Supermodularity then requires $\sigma > \eta$. By comparing the curvature in the revenue function

¹¹We derive the formulation in equation (3) from the dynamic problem of the firm in Online Appendix F.1.

to the substitutability between worker types, this condition ensures that the marginal revenue gain from rising employment of one skill type does not incentivize the firm to lower employment of another skill type. Since typical estimates of σ lie above 3 to 5, while most estimates of η lie below 2, the condition for supermodularity is compatible with standard parametrizations.

We impose Assumption (A) in the remainder of this paper. We first show that Assumption (A) guarantees existence and uniqueness of an equilibrium with a simple structure in the class of equilibria with continuous wage distributions. We then discuss how our second sufficient condition—an equilibrium refinement concept—ensures uniqueness among all possible equilibria.

Proposition 1 shows that more productive firms post higher wages.

Proposition 1. (*Wage ranking*)

Consider any equilibrium with a continuous wage offer distribution $F_s(w)$. Wages $w_s(z)$ are strictly increasing with firm productivity z . The wage function is continuous in z . The wage offer distribution satisfies $F_s(w_s(z)) = \Gamma(z)$.

Proof. See Appendix A.3. □

With Proposition 1 at hand, the distribution of workers across firms is fully determined.

Proposition 2. (*Employment distribution*)

Consider any equilibrium with a continuous wage offer distribution $F_s(w)$. The number of workers of skill s hired by firm z is given by

$$n_s(z) \equiv n_s(w_s(z)) = \frac{(1 + k_s)e_s}{MG[1 + k_s(1 - \Gamma(z))]^2}.$$

Proof. See Appendix A.4. □

Firm size in Proposition 2 depends only on the ranking of firms, $\Gamma(z)$, because firm size is fully determined by worker flows up the job ladder.¹² Building on Propositions 1 and 2, we solve explicitly for the wage distribution.

Proposition 3. (*Wage distribution*)

Consider any equilibrium with a continuous wage offer distribution $F_s(w)$. Wages are given by

$$w_s(z) = \underline{w}_s \frac{n_s(\underline{z})}{n_s(z)} + \int_{\underline{z}}^z \frac{\partial R}{\partial n_s}(z, \mathbf{n}(x)) \frac{n'_s(x) dx}{n_s(z)}.$$

Proof. See Appendix A.5. □

Proposition 3 captures the logic of the job ladder. Productive firms raise their wages to poach workers from lower-productivity firms in order to attain their target size. The equilibrium value of a worker to these lower-productivity firms is given by their marginal product of labor $\frac{\partial R}{\partial n_s}$. Competitive wage pressure for a firm with productivity z then builds up from below. Wages at a firm with productivity z are pushed up, starting from the reservation wage \underline{w}_s given in Appendix A.2, and integrating up to productivity z . The resulting wage function is a weighted average of the marginal product of

¹²We let firms choose how much recruiting effort to exert—or equivalently, how many vacancies to post—in Section 4, so that firm size also reflects the marginal product of labor.

labor at lower-ranked firms. Proposition 3 nests the wage equations with linear revenue in e.g. Burdett and Mortensen (1998) and Engbom and Moser (2021).

Proposition 3 characterizes wages having assumed that the wage offer distribution is continuous. Propositions 1 to 3 imply that the wage offer distribution that results from the choices of firms is consistent with continuity. Hence, we have guessed and verified existence of an equilibrium with a smooth wage offer distribution. As shown by Mortensen and Vishwanath (1991), wage-posting models with decreasing returns to scale can however exhibit multiple coordination equilibria due to the emergence of mass points.¹³

Our second sufficient condition is a trembling-hand equilibrium refinement concept to overcome equilibrium multiplicity. If firms make small mistakes in their wage-setting policy, no mass point can arise. When dispersion in mistakes vanishes asymptotically so that we recover the maximization problem in (3), the only equilibrium that survives is the one with a smooth wage distribution. We formalize our trembling-hand refinement in Appendix A.6 and call it Assumption (B).

Proposition 4. (*Existence and uniqueness*)

There exists a unique equilibrium among equilibria with a continuous wage offer distribution, described in Propositions 1-3. Under Assumption (B), this equilibrium is unique among all possible equilibria

Proof. See Appendix A.6. □

2 A theory of outsourcing

Having characterized the emergence of wage premia across firms in our baseline economy, we enrich our basic environment with contractor firms that provide outsourcing services and describe their impact on the economy.

2.1 Contractor firms

We introduce a measure $M_s^C \geq 0$ of identical contractor firms in each skill market s . To make the distinction clear, we now call firms that produce a consumption good ‘goods producers.’

Contractor firms hire workers in the same frictional labor markets as goods producers. They also post wages, and do so with the same recruiting technology as goods producers: every contractor firm is endowed with a unit measure of managerial time.

A given contractor firm hires in a single skill market s . Instead of producing a consumption good, contractor firms produce labor services with workers. Contractor firms sell labor services at price p_s in perfectly competitive rental markets. We endow contractors with constant returns to scale in production to make their production technology as close as possible to goods producers.¹⁴ Consistently with the comparative advantage view, contractors may be either better or worse than

¹³If a positive measure of firms coordinates on exactly the same wage, it may be optimal for other firms to post that same wage since deviating away from that mass point would imply too large a change in size given decreasing returns to production. Thus, equilibria with a smooth wage distribution may in principle co-exist with equilibria with mass points.

¹⁴Goods producers effectively transform workers in labor services in-house with constant returns, and then combine labor services to produce output with possibly decreasing returns.

goods producers at producing labor services. Contractors have comparative advantage $\tau_s \leq 1$ relative to goods producers.

Contractor thus solve the profit-maximization problem:

$$\pi_s^C = \max_w (\tau_s p_s - w) n_s(w). \quad (4)$$

We propose three micro-foundations for our measure of comparative advantage τ_s , detailed in Appendix A.7. Regardless of the micro-foundation, the comparative advantage τ_s is an exogenous parameter that captures how costly it is to outsource workers. When contractors are weakly worse than goods producers $\tau_s \leq 1$, our first micro-foundation has τ_s reflect either a productivity wedge or the inverse of an iceberg trade cost between contractor firms and goods producers. This productivity wedge or trade cost captures the idea that communication, monitoring and coordination between the goods producer and outsourced workers may be more difficult when workers are employees of another firm. As a result, some efficiency units of labor are lost. In our second micro-foundation, contractor firms combine a small amount of capital and labor according to a Cobb-Douglas production function. τ_s then simply encapsulates the price of capital. Third, $1/\tau_s$ may be interpreted as a markup charged by contractor firms. When contractors are weakly better than goods producers $\tau_s \geq 1$, τ_s reflects increasing returns to scale capturing that contractors acquire an advantage by specializing in certain labor services. Whether τ_s is above or below 1 encapsulates the net balance between those forces.

2.2 Goods producers and outsourcing

Consistently with the cost-saving view, traditional firms now face an additional possibility: rent labor services from contractors. They may still hire workers in-house in a frictional labor market. Their decision problem becomes

$$\begin{aligned} \pi(z) = & \max_{\substack{\{n_s\}_s, \{w_s\}_s, \\ \{o_s\}_s \in \{0,1\}^S}} R(z, \{n_s\}_s) - \int [(1 - o_s)w_s + o_s p_s] n_s ds \\ \text{s.t. } & n_s \leq n_s(w_s) \text{ if } o_s = 0. \end{aligned} \quad (5)$$

The indicators $o_s \in \{0, 1\}$ specify whether a goods producer outsources skill s . n_s denotes in-house labor if $o_s = 0$, and denotes outsourced labor if $o_s = 1$.¹⁵ For their problem to be well-defined, we require that R is strictly concave in \mathbf{n} .

If the goods producer hires in-house ($o_s = 0$), it effectively faces an upward-sloping labor supply curve embedded in the function of $n_s(w)$. Thus, a highly productive goods producer with a large target size n_s moves up its labor supply curve and pays high wages in-house. In contrast, if the goods producer outsources ($o_s = 1$), it faces a vertical labor supply curve at price p_s . In that case, outsourcing is more advantageous due to the upward-sloping labor supply curve $n_s(w)$: goods producers have a strong incentive to outsourcing and save on costs.

However, when a goods producer is unproductive and targets a small size n_s , it moves down its supply curve. The price of outsourcing p_s then exceeds in-house wages since it reflects both the wage

¹⁵For simplicity of exposition, we require that goods producer do not mix in-house and outsourced employment for a given skill s . None of our results are materially affected if we lift this restriction.

paid to employees of contractor firms as well as compensation to contractors. Outsourcing is then less attractive than hiring in-house. We formalize this discussion in Proposition 5 below. In what follows we always consider an equilibrium in which there is some outsourcing.

Proposition 5. (*Selection into outsourcing*)

There exists a threshold productivity \hat{z}_s , such that outsourcing of skill s occurs if and only if $z \geq \hat{z}_s$.

Proof. See Appendix A.8. □

Proposition 5 states that only the most productive goods producers in the economy outsource, precisely because they otherwise have the highest labor costs. When productivity crosses the outsourcing threshold productivity \hat{z}_s , goods producers switch from a convex to a linear cost curve. As opposed to most models of firm selection that require fixed costs (Melitz, 2003), selection into outsourcing arises because of differences in the shape of two possible variable cost functions. Switching between cost functions lets firms expand their activities, as described in Proposition 6 below.

Proposition 6. (*Productivity effect*)

Revenues and labor demand increase upon outsourcing: $R(\hat{z}_s^+, \mathbf{n}^(\hat{z}_s^+)) > R(\hat{z}_s^-, \mathbf{n}^*(\hat{z}_s^-))$ and $n_s^*(\hat{z}_s^+) > n_s^*(\hat{z}_s^-)$.*

Proof. See Appendix A.9. □

Proposition 6 encapsulates the productivity effect of outsourcing. Outsourcing benefits goods producers by letting them expand up to their preferred scale. Combined with Proposition 5, these observations imply that outsourcing effectively reallocates labor to the most productive firms in the economy. Those firms were precisely under-sized absent outsourcing due to constraints on managerial time. Outsourcing helps firms bypass those managerial constraints and improves the allocation of labor in the economy.

However, outsourcing also changes the wage structure in the labor market, with distinct implications for inequality. How the wage structure reacts to outsourcing in turn depends on whether contractors have a comparative advantage or disadvantage.

Proposition 7. (*Distributional effect*)

(a) *When contractors do not have a comparative advantage $\tau_s \leq 1$, they pay lower wages than the marginal goods producer: for any contractor wage $w_s^{\text{cont.}}$, $w_s^{\text{cont.}} < w_s(\hat{z}_s) \leq p_s$.*

If, in addition, contractors face free entry $\pi^C = 0$, they all pay the reservation wage \underline{w}_s .

(b) *Suppose that there is a single worker type and that $0 < \rho \leq -\frac{nR_{nn}}{R_n} \leq \bar{\rho} < 1$. When contractors have a high enough comparative advantage $\tau_s \gg 1$, they pay higher wages than the marginal goods producer $w_s^{\text{cont.}} > w_s(\hat{z}_s)$, with $w_s(\hat{z}_s) \leq p_s$.*

Proof. See Appendix A.10. □

To understand the forces at work in Proposition 7, consider the neutral case $\tau_s = 1$ first. In-house firms at \hat{z}_s equate the marginal product of labor to their marginal cost, that in turn encapsulates the upward-sloping labor supply curve. By contrast, outsourcing firms at \hat{z}_s equate the marginal product of labor to the price of labor services p_s on a vertical labor supply curve. Thus, outsourcing lowers the marginal product of labor of firms down to p_s . But the price of outsourcing services p_s is also the marginal product of labor for contractor firms. Hence, the marginal product of labor at contractor firms is below the marginal product of labor at in-house goods producers at \hat{z}_s .

As described in Proposition 3, wages track the marginal product of labor. Although this co-movement includes a non-constant monopsony markdown, comparing marginal products of labor is enough to compare wages. We conclude that contractor firms pay lower wages than in-house firms that would also outsource at the margin.

Next consider the strict technological disadvantage $\tau_s < 1$. Technological disadvantage only reduces the marginal product of labor of contractors relative to the marginal outsourcing firm at \hat{z}_s . This technological disadvantage further widens the gap between contractor wages and in-house marginal firms. When contractors instead have a large enough comparative advantage $\tau_s \gg 1$, their marginal product outstrips that of goods producers and so they pay higher wages.

This basic result captures the distributional effect of outsourcing. When a goods producer at \hat{z}_s decides to outsource for idiosyncratic reasons, and its workers transition to contractor firms, they experience a discrete wage change. In the weakly neutral case $\tau_s \leq 1$, this wage change is a wage penalty, as in Goldschmidt and Schmieder (2017). In the aggregate, the labor reallocation that follows a rise in outsourcing depresses earnings by moving workers away from the highest-paying goods producers and towards low-paying contractor firms.

The magnitude of earning losses depends on the exact wages paid by contractors. Although we empirically measure and structurally match the outsourcing wage penalty in the following sections, we can also make progress by imposing additional structure. A natural although perhaps strong assumption is that competition between contractors exhausts all rents—it may be easier to replicate a contracting agency than a blueprint for a particular consumption good. When this stiff competition materializes as a free-entry condition, contractors can only sustain their business by paying the lowest wage in the economy, the reservation wage \underline{w}_s . In that case, the outsourcing wage penalty is maximal.

Regardless of comparative advantage and entry, outsourcing always removes the best-paying opportunities in the labor market. Because no goods producer finds it profitable to hire a worker above the price of outsourcing, in-house wages are capped by p_s .

Our micro-foundation of monopsony power and the labor supply curve ties together wages and worker flows and delivers additional testable implications. Labor market frictions imply that workers flow towards high-paying firms. Thus, the wage ordering between contractors and goods producers from Proposition 7 translates immediately into an ordering for several other labor market statistics. We define the fraction of hires from employment of any firm paying wage w as

$$\text{HE}_s(w) = \frac{q_s(1 - \phi_s)G_s(w)}{q_s[\phi_s + (1 - \phi_s)G_s(w)]}, \quad (6)$$

where $\phi_s = \frac{\lambda_s^U u_s}{\lambda_s^U u_s + \lambda_s^E(1 - u_s)}$ is the aggregate fraction of hires from unemployment, and q_s denotes the

vacancy fill rate. The churn rate—or total separation rate—of a firm paying w is the sum of separations to unemployment and employment:

$$\text{Churn}_s(w) = \frac{[\delta_s + \lambda_s^E(1 - F_s(w))]n_s(w)}{n_s(w)}. \quad (7)$$

Churn measures how much firm w turns its workers over to maintain a stable size. We also define net poaching of a firm paying wage w as the difference between the hire rate from employment and the quit rate to employment:

$$\text{NP}_s(w) = \frac{q_s(1 - \phi_s)G_s(w) - \lambda_s^E(1 - F_s(w))n_s(w)}{n_s(w)}. \quad (8)$$

Net poaching is a commonly used revealed preference statistic to assess how attractive a firm appears to workers (Sorkin, 2018, Haltiwanger et al., 2018, Bilal et al., 2022). Proposition 8 characterizes how contractors compare to goods producers along all those labor market statistics.

Proposition 8. (*Labor market statistics of contractors*)

- (a) *When contractors do not have a comparative advantage $\tau_s \leq 1$, they hire less from employment, have higher churn and lower net poaching than the marginal goods producer: $\text{HE}_s^{\text{cont.}} < \text{HE}_s(\hat{z}_s)$, $\text{Churn}_s^{\text{cont.}} > \text{Churn}_s(\hat{z}_s)$ and $\text{NP}_s^{\text{cont.}} < \text{NP}_s(\hat{z}_s)$.*
- (b) *Suppose that there is a single worker type and that $0 < \underline{\rho} \leq -\frac{nR_{nn}}{R_n} \leq \bar{\rho} < 1$. When contractors have a high enough comparative advantage $\tau_s \gg 1$, they hire more from employment, have lower churn and higher net poaching than the marginal goods producer: $\text{HE}_s^{\text{cont.}} > \text{HE}_s(\hat{z}_s)$, $\text{Churn}_s^{\text{cont.}} < \text{Churn}_s(\hat{z}_s)$ and $\text{NP}_s^{\text{cont.}} > \text{NP}_s(\hat{z}_s)$.*

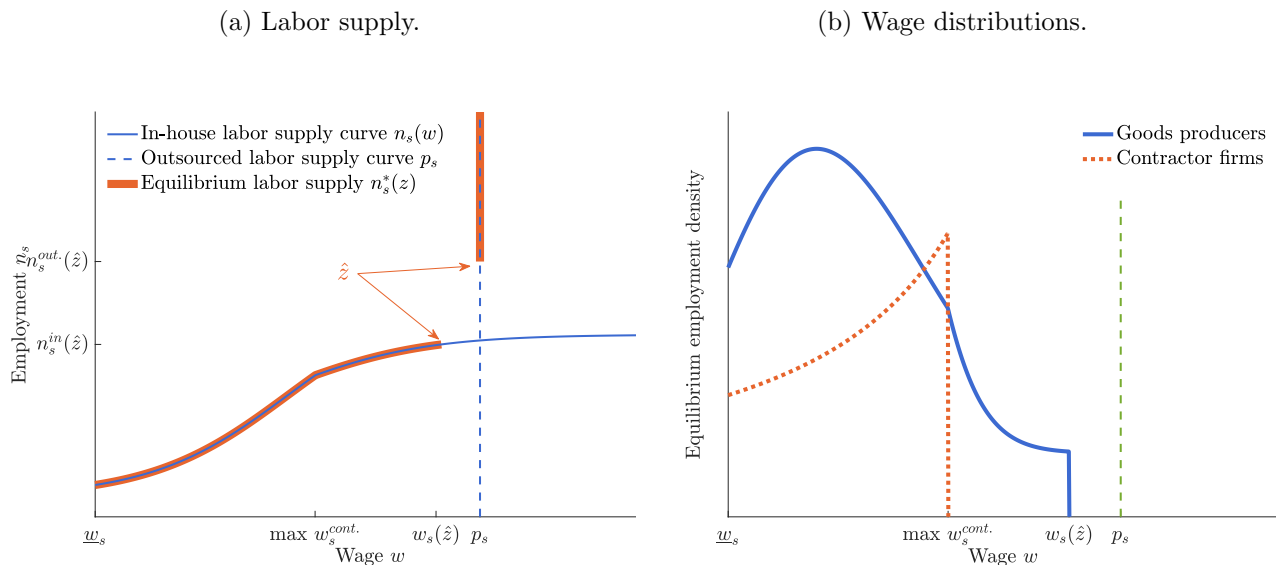
Proof. See Appendix A.11. □

Proposition 7 ranks wages of contractors and wages of the marginal goods producer, and so our labor market statistics are also ranked. In the weakly neutral case, contractors hire less from employment than the marginal firm, have higher churn, and have lower net poaching. The comparative advantage case reverses these predictions. All these comparisons are the outcome of the job ladder that emerges in equilibrium.

Finally, the ranking between goods producers and contractors on the job ladder translates into size differences. In the weakly neutral case $\tau_s \leq 1$, contractors pay less than the marginal goods producer. Because both face the same labor supply curve (2), contractors have fewer employees than the marginal goods producer. In the comparative advantage case $\tau_s \gg 1$, contractors have more employees than the marginal goods producer. We return to this comparison in Section 4, where we enrich our environment to jointly match wage and size differences between contractors and goods producers.

Together, Propositions 5 to 8 characterize the key tension between the productivity and distributional effects of outsourcing. The next section describes how these forces determine the equilibrium wage distribution.

Figure 1: Labor supply and wage distributions in equilibrium.



Note: Panel (a): in-house, outsourced and equilibrium labor supply curves of goods producers. Panel (b): equilibrium wage distributions of contractors and goods producers.

2.3 Outsourcing equilibrium

To close the description of our economy, we determine the price of outsourced labor services p_s for each skill s . It is pinned down by the market clearing condition:

$$M^G \int_{\hat{z}_s}^{\bar{z}} n_s^*(z) d\Gamma(z) = \tau_s M^C \int n_s^{\text{cont.}}(w) dF_s(w). \quad (9)$$

With some outsourcing, the equilibrium has several regions. For brevity, we describe the simplest structure of the equilibrium in this section, and provide a full description in Appendix A.12. We focus on the empirically relevant weakly neutral case $\tau_s \leq 1$ and when some goods producers and some contractors operate at the reservation wage. Figure 1 describes the equilibrium.

In the first, low wage region, goods producers operate in-house. Their labor supply coincides with the in-house supply curve as shown in Figure 1(a). In this region, goods producers compete with contractor firms as shown in Figure 1(b). The wage distribution mixes both types of firms.

The second region arises once wages reach the maximal value that contractors pay, $\max w_s^{\text{cont.}}$. Contractors no longer hire, and only highly productive goods producers who pay high wages operate. In this second region, goods producers do not compete with contractor firms. A given wage increase does not attract as many workers as in the low-wage region, and so the labor supply curve and the wage distribution may have a kink at $\max w_s^{\text{cont.}}$, as shown in panels (a) and (b).

Once goods producers become productive enough, they find outsourcing preferable to hiring in-house. This change happens at productivity \hat{z} and wage $w_s(\hat{z})$. Proposition 6 ensures that employment jumps up as goods producers switch from the upward-sloping labor supply curve to the vertical one depicted in panel (a). Proposition 7 does not restrict the size of the gap between the maximal in-house wage and the outsourcing price beyond $w_s(\hat{z}) \leq p_s$. Thus, we represent the equilibrium with a gap between both prices and quantities. When goods producers no longer operate, the employment density

drops to zero in panel (b), as there are no firms left to hire workers at higher wages than $w_s(\hat{z})$.

2.4 An empirical model of outsourcing

Not only do Propositions 5 to 7 characterize the productivity and distributional effects of outsourcing, they also provide implications that we test in Section 3 below. To connect our model with the richness of the data and emphasize possible identification challenges, we supplement our theory with additional heterogeneity.

In addition to productivity z , goods producers now face idiosyncratic outsourcing cost shocks $\boldsymbol{\varepsilon} \equiv \{\varepsilon_s\}_s$, with $1 \leq \varepsilon_s \leq \bar{\varepsilon}_s$ for all s . These shocks may be arbitrarily correlated with productivity z . They capture the idea that some managers are particularly apt at harnessing outsourcing, or that the specific production process of a given firm is well-suited for outsourcing. When deciding to hire in-house or to outsource, goods producers then solve

$$\begin{aligned} \pi(z, \boldsymbol{\varepsilon}) &= \max_{\substack{\{n_s\}_s, \{w_s\}_s, \\ \{o_s\}_s \in \{0,1\}^S}} R(z, \{n_s\}_s) - \int \left[(1 - o_s)w_s + o_s p_s \varepsilon_s \right] n_s ds & (10) \\ \text{s.t. } & n_s \leq n_s(w_s) \text{ if } o_s = 0. \end{aligned}$$

The only difference between the decision problem (10) and the decision problem (5) is that idiosyncratic outsourcing cost shocks ε_s introduce a randomization of the effective price of outsourcing across firms. As a consequence, our results remain unchanged after conditioning on outsourcing costs.

Corollary 1. (*Selection into outsourcing with cost shocks*)

There exists a threshold productivity function $\hat{z}_s(\boldsymbol{\varepsilon})$, such that outsourcing occurs if and only if $z \geq \hat{z}_s(\boldsymbol{\varepsilon})$. The threshold $\hat{z}_s(\boldsymbol{\varepsilon})$ is increasing in every ε_s . Outsourcing expenditures $\mathcal{E}(z, \boldsymbol{\varepsilon}) = \int p_s \varepsilon_s o_s^*(z, \boldsymbol{\varepsilon}) n_s^*(z, \boldsymbol{\varepsilon}) ds$ are increasing in z .

Proof. See Appendix A.13. □

Firms face the same outsourcing decision problem as in Section 2.2 conditional on outsourcing costs ε_s . Thus, they still select into outsourcing. When the firm faces higher costs ε_s , it takes a higher productivity z to reap the benefits from outsourcing, and so the outsourcing threshold \hat{z}_s is higher. Conditional on outsourcing costs, outsourcing expenditures rise with firm productivity due to positive selection. Similarly, Corollary 2 indicates that the productivity effect arises conditional on outsourcing cost shocks.

Corollary 2. (*Productivity effect with cost shocks*)

Conditional on a vector of outsourcing cost shocks $\boldsymbol{\varepsilon}$, revenues and labor demand increase upon outsourcing: $R(\hat{z}_s(\boldsymbol{\varepsilon})^+, \boldsymbol{\varepsilon}, \mathbf{n}^*(\hat{z}_s(\boldsymbol{\varepsilon})^+)) > R(\hat{z}_s(\boldsymbol{\varepsilon})^-, \boldsymbol{\varepsilon}, \mathbf{n}^*(\hat{z}_s(\boldsymbol{\varepsilon})^-))$ and $n_s^*(\hat{z}_s(\boldsymbol{\varepsilon})^+, \boldsymbol{\varepsilon}) > n_s^*(\hat{z}_s(\boldsymbol{\varepsilon})^-, \boldsymbol{\varepsilon})$.

Proof. See Appendix A.14. □

The same logic ensures that the distributional effect emerges in an economy with idiosyncratic outsourcing costs. The marginal product of labor in-house remains above the marginal product of labor when outsourcing, which in turn is higher than the marginal product of labor at contractors.

This ranking of marginal products of labor translates into a ranking of wages and other labor market statistics along the job ladder.

Corollary 3. (*Distributional effect with cost shocks*)

- (a) *When contractors do not have a comparative advantage $\tau_s \leq 1$, they pay lower wages than all marginal goods producers: for any contractor wage $w_s^{\text{cont.}}$, $w_s^{\text{cont.}} < w_s(\hat{z}_s(\boldsymbol{\epsilon}))$. Contractors also hire less from employment, have higher churn and lower net poaching than all marginal goods producers.*
- (b) *Suppose that there is a single worker type and that $0 < \underline{\rho} \leq -\frac{nR_{nn}}{R_n} \leq \bar{\rho} < 1$. When contractors have a high enough comparative advantage $\tau_s \gg 1$, they pay higher wages than all marginal goods producers $w_s^{\text{cont.}} > w_s(\hat{z}_s(\boldsymbol{\epsilon}))$. They also hire more from employment, have lower churn and higher net poaching than all marginal goods producers.*

Proof. See Appendix A.15. □

Equipped with Corollaries 1 to 3, we are in a position to connect our theory with the data.

2.5 The comparative advantage and cost-saving views

The first link between our theory and data is the outsourcing wage penalty. Our theory is uniquely suited to use the outsourcing wage penalty to disentangle the comparative advantage view from the cost-saving view.

Consistently with the comparative advantage view, our model nests the possibility that firms outsource because contractor firms have a comparative advantage in producing services of a particular type when $\tau_s > 1$, rather than to save on costs per worker. In this case, Proposition 7 indicates that contractor firms should pay higher wages than goods producers, not lower ones. This implication is at odds with the outsourcing wage penalty documented in Dube and Kaplan (2010), Goldschmidt and Schmieder (2017), Dorn et al. (2018) and Drenik et al. (2020) for the U.S., Germany and Argentina, as well as our results in Section 3.3 for France. According to the comparative advantage view, contractors should also hire more from employment, have less churn and higher net poaching. Section 3.3 indicates that all these predictions are at odds with the data.

By contrast, Proposition 7 reveals that the cost-saving view—whereby goods producers outsource simply to save on labor costs rather than leverage productive comparative advantage—is entirely consistent with the outsourcing wage penalty. It is also consistent with our results on hiring, churn and net poaching in Section 3.3. We conclude that the cost-saving view is a more likely account of outsourcing than the comparative advantage view.

We now connect our novel implications with the data.

3 Reduced-form evidence

This section lays out how we test our theory of outsourcing. We first describe our data. Second, we discuss aggregate trends in outsourcing in France. Then, we test our three main predictions:

the distributional effect, selection into outsourcing, and the productivity effect. Finally, we rule out alternative explanations for outsourcing. We provide more details in Online Appendix E.

3.1 Data

We use a combination of administrative and survey data for France between 1996 and 2007. Our first data source is the near-universe of annual tax records of French firms (*Fichier Complet Unifié de Suse*, FICUS) that report balance sheet and income statement information. We use industry codes to identify contractor firms. We observe employment, payroll, sales and purchases of intermediate inputs, from which we construct value added. However, this dataset does not detail intermediate inputs finely enough to isolate outsourcing expenditures on the buyer side.

Our second data source is a large annual firm-level survey that details purchases of intermediates at the firm level (*Enquête Annuelle d'Entreprise*, EAE). Firms report expenditures on ‘external workers.’ External workers are employees of another firm, but that fall under a contracting agreement with the surveyed firm and are at least partially under the authority of the surveyed firm. We use expenditures on external workers as our measure of expenditures on outsourced workers. Firms remain in the survey once they enter, leading to a panel structure.

Our third data source consists of employer tax records that cover labor market outcomes for French workers (*Déclaration Annuelle de Données Sociales*, DADS). We use repeated cross-sections with the universe of French workers to construct employment and wages at the firm-occupation-year level (DADS *Postes*). We also use a 4% representative panel to study the wage penalty of outsourcing (DADS *Panel*).

Our fourth data source are customs records for the universe of trade transactions (*Donnés de Douanes*). We observe exports at the product-country-firm-year level. We use this data to construct export demand shocks at the firm-level and exploit variation in firm scale.

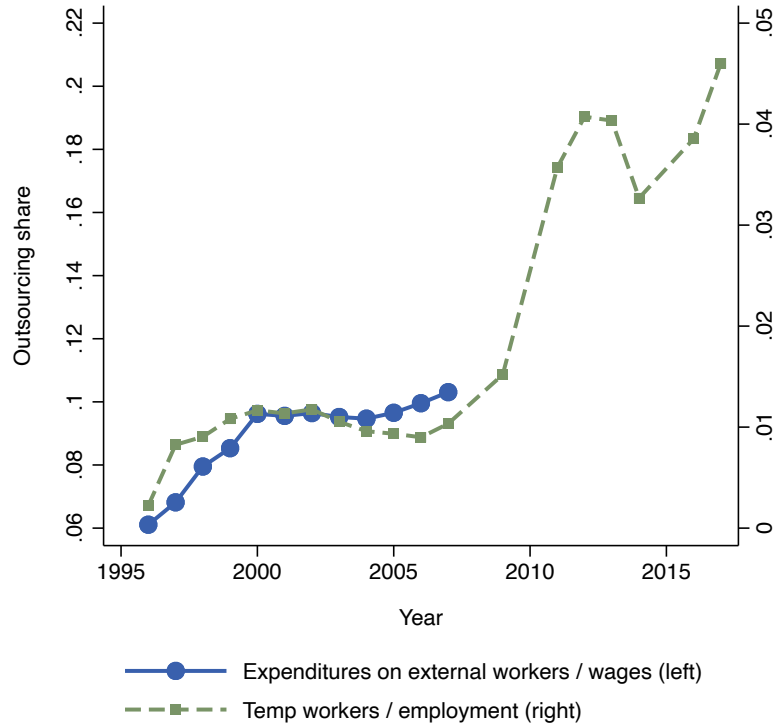
We link these four data sources together using a common firm tax identifier. For our main empirical exercises at the firm level, we aggregate years into three periods 1997-1999, 2000-2002, 2003-2007 and keep only firms with at least ten in-house employees to limit measurement error in outsourcing expenditures. We stop our main analysis in 2007 because of a large change in classification that prevents us from reliably measuring outsourcing expenditures directly in subsequent years. Our final sample consists of 173,547 firm-periods.¹⁶

3.2 Aggregate trends in outsourcing

We start by asking by how much did outsourcing rise in France. Figure 2 shows that outsourcing expenditures as a fraction of the aggregate wage bill almost double in the decade that we study: they increase from 6% in 1996 to over 10% in 2007. To infer whether the upward trend in outsourcing continues past 2007, we use a subcategory of outsourcing that we can reliably measure past 2008: temporary workers. Of course, extrapolating the trend in overall outsourcing expenditures using the

¹⁶France harmonized its data collection procedures and classification in 2008 with the rest of the European Union. The change in industry codes prevents us from identifying contractor firms after 2008. In addition, the EAE survey was discontinued and replaced by the *Enquête Sectorielle Annuelle* (ESA). The ESA includes questions about outsourcing but the response rate is substantially lower than in the EAE for firm-level expenditures, leading to severe measurement difficulties.

Figure 2: Outsourcing expenditure share and temporary work share.



Note: Solid blue line: aggregate outsourcing share, computed as aggregate expenditures on external workers relative to the aggregate in-house wage bill. Dashed green line: temporary work share, computed as total temporary workers at firms with more than 100 in-house employees, relative to all employment at firms with more than 100 in-house employees.

more restrictive group of temporary workers requires strong proportionality assumptions. But to the extent that outsourcing expenditures track the rise in the use of temporary work, Figure 2 reveals that outsourcing may account between 10% and 20% of the aggregate wage bill in France by 2016.¹⁷

This rise in aggregate outsourcing expenditures also translates into an increase in the employment share of contractor firms in the aggregate. To complement our expenditure-side data, we follow Goldschmidt and Schmieder (2017) and rely on industry and occupation codes to detect contractor firms and service workers at contractor firms. Figure 3(a) shows that the fraction of low skill workers employed at contractor firms rises from 5% in 1996 to 9% in 2007. This 4 p.p. increase in the outsourcing employment share coincides almost exactly with the increase in the outsourcing expenditure share. Figure 3(b) also reveals that the increase in the employment outsourcing share is driven specifically by service workers reallocating towards contractor firms over time. The fraction of low skill service workers employed at contractors rises by 16 p.p. over the decade we study.¹⁸

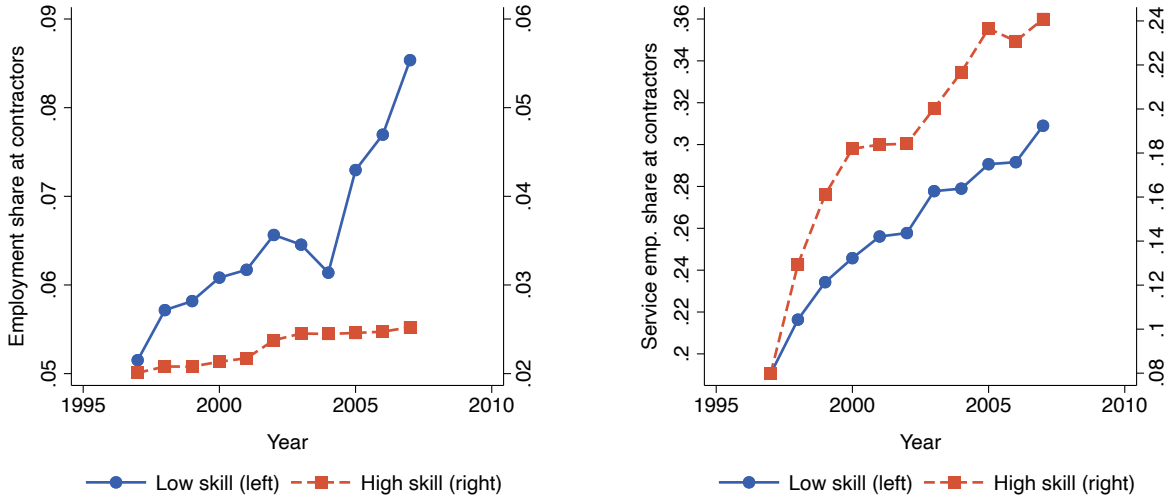
¹⁷That outsourcing represents 20% of France’s wage bill by 2016 is likely an overestimate because two regulatory changes directly impact temporary work relative to general outsourcing. In 2005 temporary work agencies are allowed to also help their workers transition into permanent work contracts at their clients. In 2009 governmental institutions are allowed to hire temporary workers.

¹⁸Our employment-based and expenditure-based measures complement each other. Our employment-based measure may miss any firm that is a contractor firm, but does not fall into our specific industry codes. Our measure using outsourcing expenditures is not subject to this limitation.

Figure 3: Outsourcing employment share.

(a) Fraction of employment at contractors.

(b) Fraction of service employment at contractors.



Note: Low skill contractor firm defined by industry codes specifically labeling firms as providing food, security, cleaning or general administrative services to other firms. High skill contractor firm defined by industry codes specifically labeling firms as providing accounting, law or consulting services to other firms.

Next, we compare salient features of firms that rely on outsourcing to firms that do not. Consistently with Propositions 5 and 6, Table 4 in Appendix B.1 shows that firms that outsource are larger, sell more and have higher value added than firms that do not outsource. Table 4 in Appendix B.1 also provides examples of 3-digit industries that outsource most and least in France. The industry that outsources the most in France is Business supplies and equipment trade. Firms in this industry mostly place orders on behalf of their client companies and outsource delivery and installation. Similarly, the Telecommunications industry manages the grid but outsources a substantial fraction of its maintenance and installation. By contrast, Manufacturing of terracota ceramics requires specific knowledge and thus does not outsource workers. Similarly, Transportation of goods and individuals into space requires highly specialized knowledge and is subject to strict security measures, and thus does not rely on outsourced workers.

3.3 The distributional effect

We start by testing the core predictions of our theory relating to the distributional effect of outsourcing. We propose two exercises following our results in Corollary 3. Our first step is to verify that there is indeed an upward-sloping labor supply curve for service workers—the premise upon which the entire distributional effect is built. Second, we show that contractor firms locate at the bottom of this job ladder: they pay lower wages, hire less from employment, have higher churn and lower net poaching.¹⁹

The upward-sloping labor supply curve. We start by verifying that larger firms pay service workers more and thus locate at the top of the job ladder. Following the large literature establishing

¹⁹Importantly, health and retirement benefits are not tied to employers in France as they are in the U.S., and so wages are the main dimension of worker compensation in our context.

the existence of a firm wage size premium, we leverage our panel data to project log wages at the worker level on the size of their employer. We control for worker fixed effects to absorb worker-level heterogeneity that may be correlated with firm scale. Following Goldschmidt and Schmieder (2017), we restrict attention to wages of workers who are in service occupations: food, security, cleaning or administrative services. We run a fixed-effect regression of the form:

$$\log w_{i,t} = \varphi_i + \phi \log \text{Total employment}_{J(i,t)} + v_{i,t}. \quad (11)$$

i indexes workers, $J(i, t)$ the employer of worker i in quarter t , and $v_{i,t}$ is a mean-zero residual. $\log w_{i,t}$ denotes the log wage, and φ_i is a worker fixed effect.

We find that larger firms indeed pay their service workers more after controlling for worker fixed effects φ_i . We display our estimates in Table 8 in Appendix B.3. Firms with 1,000 employees pay wages that are on average 9.6% higher than firms with 10 employees. We also report that wages rise with firm value added and with the overall average wage at the firm, including non-service workers. Together, these results point to an upward-sloping labor supply curve for service workers consistent with our theory.

Contractors on the job ladder. Second, we test whether contractor firms locate at the bottom of the job ladder. We identify contractor firms using the industry codes as in Section 3.2. Our goal is to measure the wage premium paid by contractors controlling for unobserved heterogeneity at the worker level. We do so with a two-way fixed effects regression as in Abowd et al. (1999), henceforth AKM:

$$\log w_{i,t} = \varphi_i + \psi_{J(i,t)} + \eta_{i,t}. \quad (12)$$

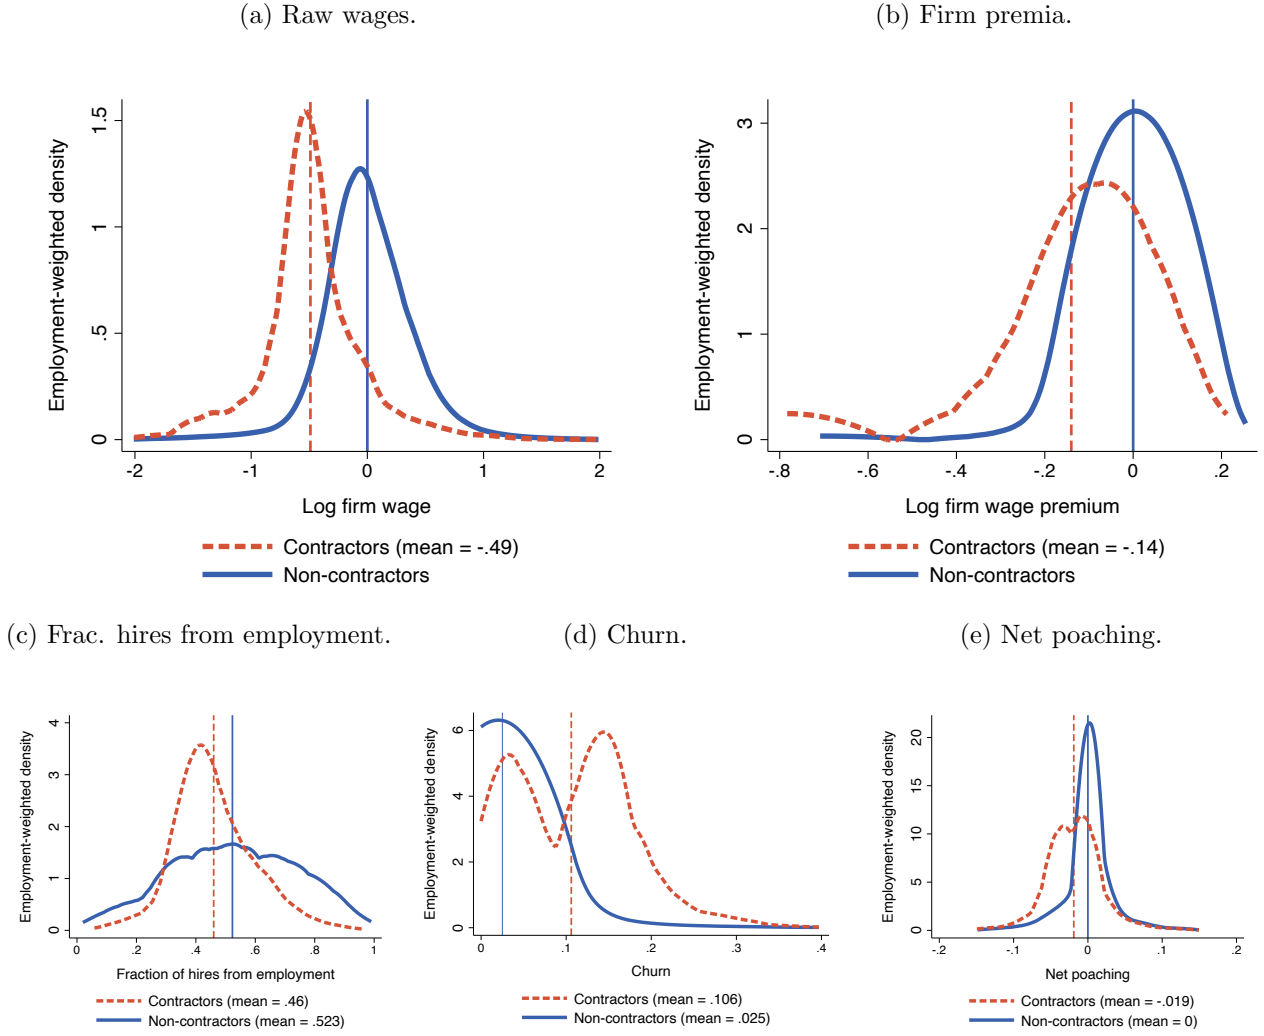
where $\psi_{J(i,t)}$ a firm fixed effect. Workers who move between firms identify separately worker and firm fixed effects if worker mobility is conditionally random (Card et al., 2013).

To limit well-known econometric difficulties linked to limited mobility bias, we follow Bonhomme et al. (2019) and group workers and firms each in 50 equally populated groups based on the unconditional mean worker and mean firm wage. We then estimate equation (12) with OLS at the group level. Our results are virtually identical when varying the number of groups between 10 and 200.

Without controlling for worker composition, contractors pay wages that are almost 50% below wages of non-contractor firms. Figure 4(a) displays the distribution of average wages at contractors and non-contractors with a kernel density plot. We show that controlling for worker composition is key to not overstating the contractor wage penalty in Figure 4(b). The penalty drops to 14% on average after removing worker fixed effects with specification (12). We also estimate the standard deviation of firm effects to be 0.14. Thus, contractors pay on average one standard deviation below the average non-contractor firm. Our results are consistent with Goldschmidt and Schmieder (2017), Dorn et al. (2018) and Drenik et al. (2020) who also find a substantial outsourcing wage penalty driven by the loss of firm wage premia.

Guided by our theory, we also propose novel measures indicating that contractors rank towards the bottom of the job ladder along other key labor market statistics. Contractor hire less from employment than goods producers as we show in Figure 4(c). The fraction of hires from employment is 6 p.p. lower at contractors. Figure 4(d) indicates that churn at contractors is 8 p.p. above that at other firms.

Figure 4: Contractors and non-contractors on the job ladder.



Note: Kernel density plots of labor market statistics of contractor firms (dashed orange) vs. non-contractor firms (solid blue). Vertical lines depict means. Raw wages: firm-level mean wages. Firm premia: firm fixed effect from the AKM regression (12). Fraction hires from employment: hires from employment divided by total hires as in equation (6). Churn: total separations divided by employment as in (7). Net poaching: hires from employment net of quits to employment divided by total employment as in (8).

Finally, contractors tend to lose their workers to other firms as shown in Figure 4(e). Net poaching at contractors is 2 p.p. below net poaching at non-contractors. This difference is substantial and corresponds to net poaching differences between young and old firms in the U.S. (Bilal et al., 2022). Together, these results indicate that contractors indeed locate at the bottom of the job ladder.

3.4 Selection into outsourcing

Having established the outsourcing wage penalty in France, we test the novel, firm-level predictions of our theory of outsourcing that links workers and firms in equilibrium. We start by testing selection into outsourcing. Corollary 1 indicates that when productivity z rises conditional on outsourcing costs $\{\varepsilon_s\}_s$, firms spend relatively more on outsourcing: $\mathcal{E}^*(z, \varepsilon)$ increases. To contrast our theory with the data, we must make three measurement choices and address several identification challenges.

First, what is productivity z ? From our model’s perspective, revenue productivity is the driver of outsourcing decisions rather than physical productivity. Still, we do not observe revenue productivity. We leverage that our model ties productivity z with revenues $R^*(z, \{\varepsilon_s\}_s)$: more productive firms sell more. To test selection into outsourcing, we thus use revenues as our main independent variable.

Second, what are revenues? Our model does not have intermediate inputs other than outsourcing. Thus, the empirical counterpart of revenues in our model is value added complemented with outsourcing expenditures which we measure in our data. We define revenues of firm f in time period t as value added VA_{ft} plus outsourcing expenditures measured as expenditures on external workers \mathcal{E}_{ft} , $R_{ft} = VA_{ft} + \mathcal{E}_{ft}$. For brevity and when unambiguous, we will sometimes refer to our measure of revenues simply as value added. We replicate all our results using instead value added VA_{ft} in Appendix B.1 and show that they remain virtually unchanged.

Third, what are outsourcing expenditures? We use log expenditures on external workers $\log \mathcal{E}_{ft}$ as our main measure of outsourcing. We also use the outsourcing share S_{ft} of firm f in time period t , defined as its expenditures on external workers \mathcal{E}_{ft} divided by the sum of its expenditures on labor $\mathcal{W}_{ft} + \mathcal{E}_{ft}$, where \mathcal{W}_{ft} denotes gross payroll: $S_{ft} = \frac{\mathcal{E}_{ft}}{\mathcal{W}_{ft} + \mathcal{E}_{ft}}$.

Log expenditures connect directly to Corollary 1, which predicts that $\frac{\partial \log \mathcal{E}}{\partial z} > 0$. However, there may be reasons beyond our model that drive a mechanical correlation between firm scale and outsourcing expenditures. Therefore, we also use the outsourcing share as an alternative dependent variable. Using the outsourcing share is a more demanding test of our theory as our model does not guarantee that the outsourcing share should always be increasing in productivity.²⁰

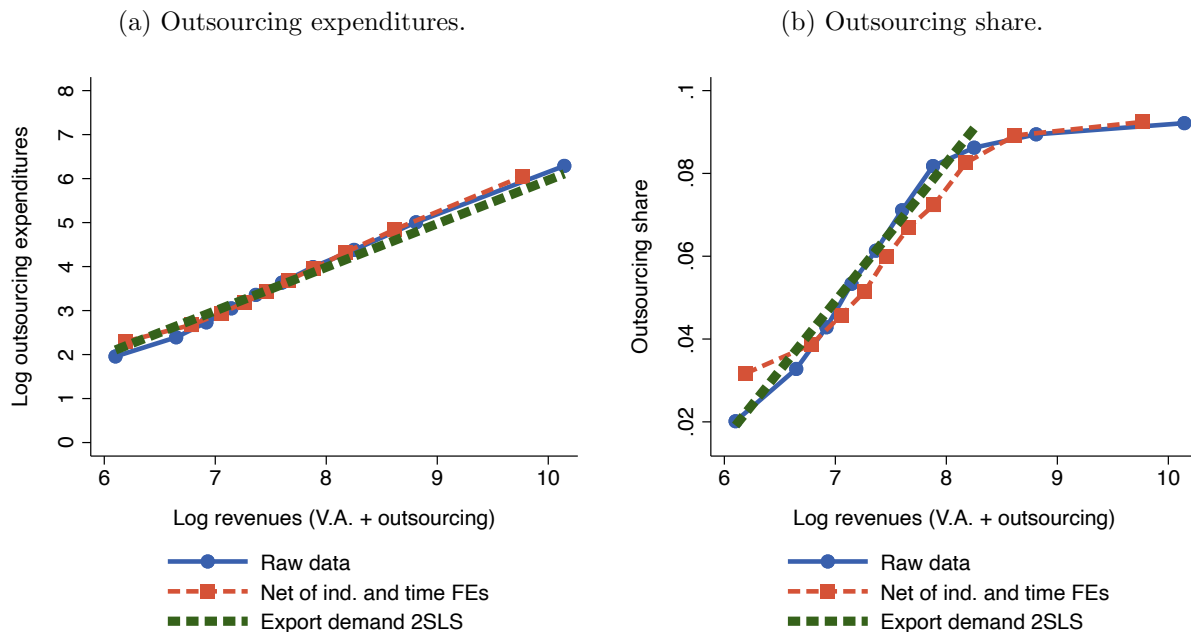
Having established our measurement strategy, we turn to our identification strategy. Selection into outsourcing translates as $\frac{\partial \log \mathcal{E}(z, \boldsymbol{\varepsilon})}{\partial z} > 0$ and, under our stricter test, $\frac{\partial S(z, \boldsymbol{\varepsilon})}{\partial z} > 0$. Since we do not measure revenue productivity directly, we seek to identify whether $\frac{\partial \log \mathcal{E}(z, \boldsymbol{\varepsilon})}{\partial z} / \frac{\partial \log R(z, \boldsymbol{\varepsilon})}{\partial z} > 0$ and $\frac{\partial S(z, \boldsymbol{\varepsilon})}{\partial z} / \frac{\partial \log R(z, \boldsymbol{\varepsilon})}{\partial z} > 0$. Identification crucially depends on the co-movement between productivity z and outsourcing costs ε_s across firms. The starkest identification assumption is that there is no dispersion in outsourcing costs ε_s across firms. In this case, we may simply correlate revenues $\log R_{ft}$ with outsourcing expenditures $\log \mathcal{E}_{ft}$ or the outsourcing share S_{ft} across firms to test whether there is selection into outsourcing.

Figure 5 plots outsourcing by decile of revenues. Figure 5(a) reports how outsourcing expenditures change with revenues. Consistently with Table 4, Figure 5 reveals that in the raw data, high revenue firms outsource more. The relationship is nearly log-linear for outsourcing expenditures. Figure 5(b) reports the outsourcing share by revenues. A firm in the first decile of revenue spends 2% of its labor costs on outsourced labor, while a firm in the tenth decile of revenue spends over 8%. This upward-sloping relationship is not an artifact of industry composition or time trends, as shown by the residualized relationship. Under the assumption of no outsourcing costs ε_s , the slopes in both panels of Figure 5 support the view that firms select into outsourcing.

Of course, the relationship depicted in Figure 5 could be the result of variation in outsourcing

²⁰For instance, if the revenue function is non-homothetic and firms increasingly rely on high skill workers as they expand, in-house payroll can rise faster than outsourcing expenditures. In that case, the outsourcing share decreases with productivity. Figure 17 in Appendix B.1 suggests that skill deepening may indeed be affecting the outsourcing share at the far end of the productivity distribution, although our estimates are noisy.

Figure 5: Outsourcing by revenues.



Note: Solid blue line: raw data. Dashed orange line: after removing 3-digit industry and time period fixed effects from log outsourcing expenditures, the outsourcing share and log value added. Thick dashed green line: 2SLS estimate using the export demand shift-share instrument in equation (14). Panel (a): log outsourcing expenditures. Panel (b): outsourcing share.

costs across firms. Perhaps the particular production process of some firms lends itself well to outsourcing, while other firms find it more difficult to outsource. Since revenues $R(z, \varepsilon)$, outsourcing expenditures $\mathcal{E}(z, \varepsilon)$ and the outsourcing share $S(z, \varepsilon)$ are all likely decreasing in ε_s , the unconditional correlation between both may be positive even if outsourcing expenditures and the outsourcing share are in fact independent from productivity z . In that case, selection into outsourcing based on firm productivity would be spuriously identified. This threat to identification is particularly problematic when productivity z and outsourcing costs ε_s are correlated across firms. For instance, if the CEO of a firm is particularly skilled at harnessing the advantages of information technology, it may improve overall productivity at the firm at the same time as making outsourcing easier for reasons unrelated to selection into outsourcing.

To circumvent this difficulty we focus next on the conditional correlation between revenues R_{ft} , outsourcing expenditures \mathcal{E}_{ft} and the outsourcing share S_{ft} within firms. This conditional correlation identifies selection into outsourcing if outsourcing costs are constant within firms. Under this assumption, productivity z and outsourcing costs ε_s may be arbitrarily correlated across firms in levels, but not in changes. For instance, a CEO may be particularly skilled at reorganizing its firm to outsource, but their ability may not vary over time. Up to a first-order approximation, the linear regression

$$Y_{ft} = \alpha_t + \beta_f + \gamma \log R_{ft} + \eta_{ft}, \quad Y_{ft} \in \{\log \mathcal{E}_{ft}, S_{ft}\}, \quad (13)$$

then identifies selection into outsourcing. α_t is a time period fixed effect, β_f a firm fixed effect, and η_{ft} a mean zero residual. The firm fixed effect β_f absorbs the co-movement between long-run productivity and outsourcing costs across firms. The coefficient γ captures selection into outsourcing. It is identified by within-firm changes in revenues and outsourcing driven by shocks to revenue productivity.

We report the regression version of our results without and with firm fixed effects in Tables 5 and 6 in Appendix B.1. Table 5 displays our results for outsourcing expenditures, and Table 6 reports our results for the outsourcing share. In both tables, columns (1-2) show the regression analog of Figure 5. A 10 log points increase in revenues is associated with a 11 log points increase in log outsourcing expenditures and a 0.18 p.p. increase in the outsourcing share. The estimate for outsourcing expenditures slightly shrinks to 8 log points when we focus on within-firm variation in column (3) by including firm fixed effects. For the outsourcing share, the estimate rises to 0.26 p.p. within firms.

An additional threat to identification arises when outsourcing costs vary over time at the firm level. In that case, the estimate of γ in the OLS regression (13) may simply capture the joint dependence of revenues and the outsourcing share on outsourcing costs. This threat is even more problematic when changes in outsourcing costs are also correlated with changes in revenue productivity. Returning to our example, the ability of the CEO to improve productivity and make outsourcing easier may vary over time. In this case, changes in productivity and outsourcing costs are correlated even within firm, and the OLS regression (13) is inadequate to identify selection into outsourcing.

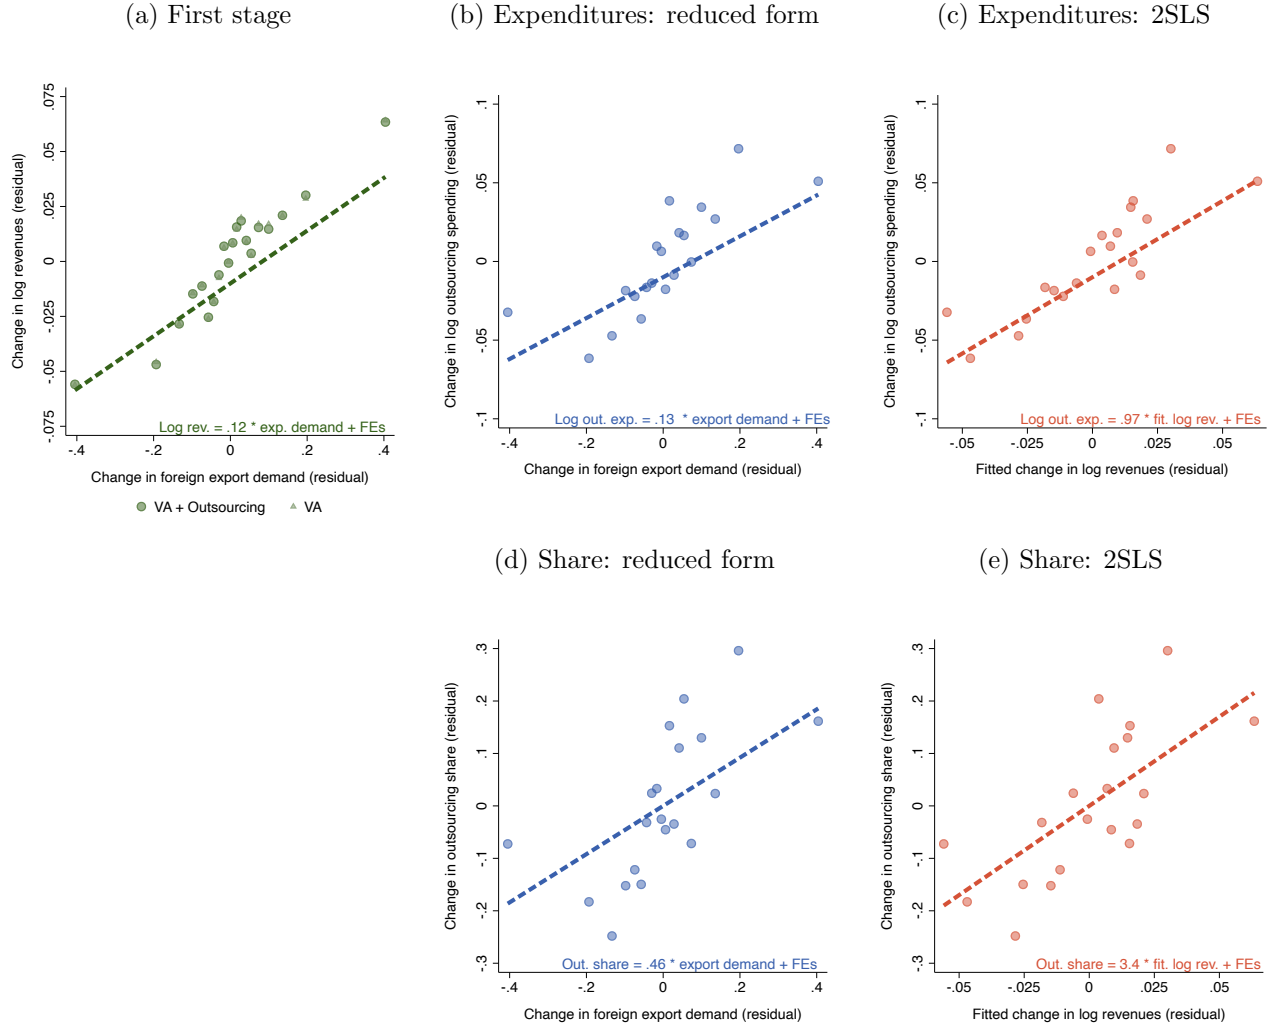
We turn to an instrumental variable strategy to address this identification challenge. We identify shocks to revenue productivity z that are plausibly unrelated to outsourcing cost shocks ε_s . We consider foreign export demand shocks. We exploit the granularity of our customs data and follow Hummels et al. (2014). We first construct firm-level export shares in the first time period, $\pi_{f,t_0,j}$, across 4-digit industry-country pairs j . We then interact those shares with export demand growth $\Delta \log X_{j,t,-f}$ in industry-country pair j between time periods t_0 and t , excluding firm f 's exports. The instrument $Z_{f,t}$ for revenues R_{ft} is thus defined as a shift-share:

$$Z_{f,t} = \sum_j \pi_{f,t_0,j} \Delta \log X_{j,t,-f}. \quad (14)$$

The identifying variation in the instrument follows from changes in foreign export demand. Consider firm f that exports luxury handbags to South Korea in the initial period. If South Korean demand for luxury handbags subsequently grows, firm f will face an increase in demand. Our exclusion restriction is that this rise in firm f 's demand that follows from South Korea's higher demand for luxury handbags is unrelated to firm f 's own ability to outsource (Borusyak et al., 2021). In that case, foreign export demand shocks raise revenue productivity and $Z_{f,t}$ is a valid instrument for changes in firm revenues. We display the distribution of shares $\pi_{f,t_0,j}$, changes in export demand $\Delta \log X_{j,t,-f}$ and the resulting instrument $Z_{f,t}$ in Figure 16, Appendix B.1. Up to a first-order approximation, the 2SLS estimate of γ in equation (13) then captures the ratio of average partial derivatives $\mathbb{E} \left[\frac{\partial \log \mathcal{E}(z,\boldsymbol{\varepsilon})}{\partial z} \right] / \mathbb{E} \left[\frac{\partial \log R(z,\boldsymbol{\varepsilon})}{\partial z} \right] > 0$ and $\mathbb{E} \left[\frac{\partial S(z,\boldsymbol{\varepsilon})}{\partial z} \right] / \mathbb{E} \left[\frac{\partial \log R(z,\boldsymbol{\varepsilon})}{\partial z} \right] > 0$, and thus selection into outsourcing.

The main concern for identification under our instrumental variable strategy arises if firms that happen to have large export exposure to markets that grow are also firms who experience systematic reductions in their costs of outsourcing. This type of correlation would arise if managers are particularly foresighted and invest ahead of time in export destinations that will grow in the future, while subsequently restructuring the firm to make it more amenable to outsourcing. While possible in principle, the systematic prevalence of such situations is unlikely in practice given the large fixed costs typically associated with exporting to new markets

Figure 6: Selection into outsourcing: Instrumental variable approach.



Note: Bin-scatterplot of the first stage 6(a), reduced form 6(b) and two stage least square 6(c) estimates for selection into outsourcing with log outsourcing expenditures as dependent variable. Bin-scatterplot of the reduced form 6(d) and two stage least square 6(e) estimates for selection into outsourcing with the outsourcing share as dependent variable. Panels include regression lines for 20 bins. Coefficients may differ from full sample regression coefficients reported in Tables 5 and 6, Appendix B.1.

Figure 6 represents our instrumental variable strategy graphically. Panel 6(a) shows a strong first stage: export demand growth predicts value added growth, with an F-statistic of 332 well above conventional thresholds for weak instruments (Stock and Yogo, 2005). Panels 6(b) and 6(d) reveal that export demand growth also leads to growth in firm-level outsourcing expenditures and outsourcing shares. As a result, panels 6(c) and 6(e) uncover a positive relationship between revenue productivity and outsourcing.

We collect our estimates using our full firm-level sample in Tables 5 and 6, Appendix B.1. A 10% increase in value added driven by export demand shocks leads outsourcing expenditure to rise by 10% and the outsourcing share to rise by 0.33 p.p.²¹ All our estimates are economically and statistically significant at the 0.1% level. Since our instrument only affects exporters, we also confirm that exporters

²¹In panel 6(e), the coefficient is 0.34 p.p. when we run the regression at the bin level rather than at the firm level.

exhibit a similar within-firm OLS relationship between revenues and outsourcing relative to all firms in column (5) in Tables 5 and 6, Appendix B.1. The thick dashed green lines in Figure 5 depict our 2SLS estimates graphically and show that they are quantitatively comparable to the slope of the cross-sectional relationship.

We verify the robustness of our results with other metrics of firm performance. Figure 15 in Appendix B.1 together with column (6) of Tables 5 and 6 indicate that our OLS and 2SLS estimates are virtually identical when we use value added without including outsourcing expenditures as our main measure of revenues. We also use employment and value added per worker as alternative measures of firm performance. Columns (7-8) of Tables 5 and 6 in Appendix B.1 all point to an economically meaningful and statistically significant effect of firm scale on outsourcing regardless of the particular metric we use. We conclude that firms select into outsourcing.²²

3.5 The productivity effect

We now test the third core prediction of our theory: the productivity effect of outsourcing. Corollary 2 indicates that when outsourcing costs ε_s decline so much that firms outsource, they expand: $R(z, \varepsilon)$ rises. To confront this prediction with the data, we face similar measurement and identification challenges to those described in Section 3.4.

We use similar variable definitions as in the previous section. Our main dependent variable is value added at the firm and period level, VA_{ft} . The independent variable is the outsourcing share S_{ft} . We seek to identify whether $\frac{\partial \log R(z, \varepsilon)}{\partial (1/\varepsilon_s)} > 0$. Since we do not measure outsourcing costs directly, we ask whether $\frac{\partial \log R(z, \varepsilon)}{\partial (1/\varepsilon_s)} / \frac{\partial S(z, \varepsilon)}{\partial (1/\varepsilon_s)} > 0$.

As in Section 3.4, we could use a conditional correlation between revenues and the outsourcing share to identify the productivity effect if revenue productivity z was constant within firms. In that case, we could run a regression

$$\log VA_{ft} = \alpha'_t + \beta'_f + \gamma' S_{ft} + \varepsilon'_{ft} \quad (15)$$

and use the OLS estimate of γ' to identify the productivity effect. Of course, the identification assumption we need is inconsistent with those in Section 3.4 and is unlikely to hold. In addition, changes in outsourcing costs may be correlated with changes in revenue productivity within firm. The same CEO example illustrates the main threat to identification. The ability of a CEO to improve productivity and make outsourcing easier may vary over time. Finally, the outsourcing share is likely to contain more measurement error than value added at the firm level. Measurement error would bias the OLS estimate of γ' towards zero.²³

We implement an instrumental variable strategy to address these identification challenges. We leverage that firms are differentially exposed to service occupations o : food, security, cleaning or

²²Abraham and Taylor (1996) find that establishments with larger employment are less likely to outsource. Several factors may explain the difference between our findings and theirs. First, they consider establishment-level data, not firm-level data. Second, the minimum size for establishments to be included in their sample ranges between 20 and 100 employees. Figure 20(b) reveals that most of the relevant variation is concentrated below 50 employees. Third, they use cross-sectional survey data with 2,700 establishments, while we have administrative panel data with 173,547 firm-periods.

²³We obtain value added from the FICUS data which is based on administrative tax returns. By contrast, outsourcing expenditures come from the EAE survey and correspond to much smaller income statement categories.

general administrative occupations. Our first step is to construct average outsourcing expenditures on occupation o , denoted by $\Omega_{o,t,-f}$. We interact initial payroll shares ω_{f,o,t_0} with firm-level outsourcing expenditures $\mathcal{E}_{f,t}$. We then sum across firms to obtain $\Omega_{o,t,-f} = \frac{1}{N-f} \sum_{f' \neq f} \omega_{f',o,t_0} \mathcal{E}_{f',t}$.

Next, we construct a predicted outsourcing share for firm f by interacting initial payroll shares of firm f with average outsourcing expenditures in the same occupations $\Omega_{o,t,-f}$:

$$\widehat{S}_{f,t} = \frac{\sum_o \omega_{f,o,t_0} \Omega_{o,t,-f}}{\mathcal{W}_{f,t_0} + \mathcal{E}_{f,t_0}}.$$

We then instrument the change in the outsourcing share of any firm using changes in the predicted outsourcing share:²⁴

$$Z'_{f,t} = \Delta \widehat{S}_{f,t}$$

The initial exposure of firms to different occupations generates the identifying variation in the instrument. For instance, firm f that produces luxury handbags also needs to hire many security guards in-house in the initial period to secure its warehouses. Over time, average outsourcing expenditures on security guards are rising, revealing economy-wide declines in outsourcing costs specifically for security guards. Our instrument infers that firm f is particularly exposed to these costs declines, and thus should experience a substantial rise in its outsourcing share. We interpret this differential exposure as idiosyncratic changes in outsourcing costs ε_s .

Our exclusion restriction is that the resulting decline in idiosyncratic outsourcing costs is unrelated to changes in revenue productivity (Goldsmith-Pinkham et al., 2020). We display the distribution of shares ω_{f,o,t_0} , changes in outsourcing expenditures $\Delta \Omega_{o,t,-f}$ and the resulting instrument $Z'_{f,t}$ in Figure 18 in Appendix B.2. Up to a first-order approximation, the 2SLS estimate of γ' in (15) identifies the ratio of average partial derivatives $\mathbb{E} \left[\frac{\partial R(z,\varepsilon)}{\partial(1/\varepsilon_s)} \right] / \mathbb{E} \left[\frac{\partial S(z,\varepsilon)}{\partial(1/\varepsilon_s)} \right]$ and thus the productivity effect of outsourcing.

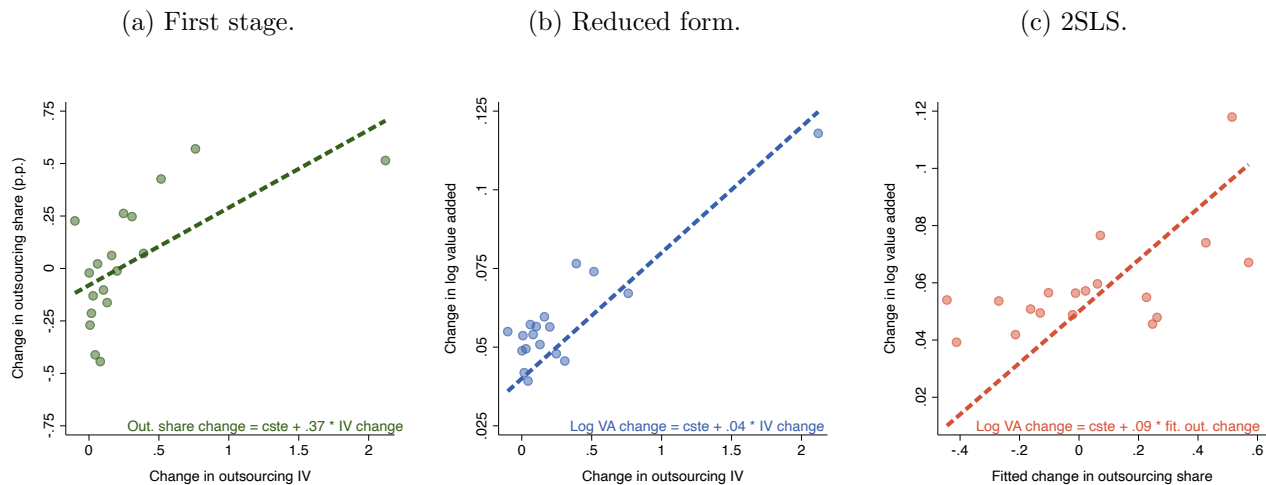
The main threat to identification arises if firms with many security guards in the initial period are also those that experience subsequent productivity growth. Once more, a particularly foresighted manager may have chosen a production structure that heavily relies on security guards, anticipating future outsourcing cost declines. If this manager is also particularly talented and makes its firm expand for reasons unrelated to outsourcing, our exclusion restriction is violated. While the pervasive adoption of such clear-sighted managerial practices would certainly be beneficial to the French economy, predicting which particular occupations will experience outsourcing cost declines in the future and distorting the production process of the firm towards them ex-ante is a high bar.

Figure 7 represents our instrumental variable strategy graphically. Panel 7(a) reveals a strong and positive first stage. Panel 7(b) reveals that growth in the predicted outsourcing share also leads to growth in firm-level value added. As a result, panel 7(c) uncovers a positive relationship between outsourcing and value added.

We collect our estimates using our full firm-level sample in Table 7, Appendix B.2. The first stage F-statistic is 23.8. A 1 p.p. increase in the outsourcing share driven by outsourcing cost shocks leads

²⁴The first stage associated with our instrument is positive when $\mathbb{E} \left[\frac{\partial S(z,\varepsilon)}{\partial(1/\varepsilon_s)} \right] > 0$, i.e. if a decline in outsourcing costs lead firms to spend more on outsourcing. We show that this is the case in Appendix A.16 for the setup from Section 4. Ultimately however, the sign of the first stage is an empirical question.

Figure 7: Productivity effect: Instrumental variable approach.



Note: Bin-scatterplot of the first stage 7(a), reduced form 7(b) and two stage least square 7(c) estimates for selection into outsourcing. Panels include regression lines for 20 bins. Coefficients may differ from full sample regression coefficients reported in Table 7, Appendix B.2.

value added to rise by 7.9%. This point estimate is economically and statistically significant at the 0.1% level. We conclude that outsourcing has a positive productivity effect at the firm level: firms that outsource produce more.²⁵

3.6 Alternative explanations

In principle, mechanisms that we do not emphasize through our theory may also lead firms to outsource. In this final subsection, we rule out three prominent alternative explanations as key drivers of outsourcing in France.

Volatility. The first explanation is that firms outsource because they value the associated flexibility when demand is volatile. This explanation conflicts with the data along two dimensions. First, we should observe a negative relationship between firm scale and outsourcing since small firms are more volatile. Figure 19(a) in Appendix B.4 indicates precisely the opposite. Second, Figure 19(b) in Appendix B.4 shows that industries with higher value added volatility actually rely less on outsourcing, not more. At the firm level however, we find some evidence in favor of a modest relationship between volatility and outsourcing, consistently with Abraham and Taylor (1996). Table 9 in Appendix B.4 reveals that firm-level volatility is positively associated with outsourcing. Importantly however, the standardized coefficient is two to three times smaller than for firm scale. Thus, volatility matters substantially less than firm scale. We conclude that the primary reason that firms outsource is unlikely to be workforce flexibility.

Equity. Second, the upward-sloping labor supply curve faced by firms could be partly due to equity concerns rather than scarce managerial time. Our theory applies equally well if the labor supply curve

²⁵We cannot test directly the productivity for labor demand because we only observe outsourcing expenditures rather than quantities at the firm level.

is partly generated by equity concerns, but we still investigate whether equity concerns per se lead to outsourcing.

The equity motive for outsourcing implies that firms with more unequal pay structures then have the strongest incentives to outsource. This mechanism turns out to be at odds with the data. Table 10 in Appendix B.4 indicates that, if anything, firms with more unequal pay structures outsource less, not more. Our results hold when within-firm inequality is measured as the standard deviation of log wages or the 90th to 10th percentile ratio. They also hold conditional on log revenues and in changes over time.

Our results are consistent with the growing literature that documents equity concerns (e.g. Card et al., 2012, Breza et al., 2017). Indeed, this strand of work finds that equity concerns are primarily binding across workers within the same occupation, rather than across occupations or worker performance categories. We conclude that equity concerns are not a dominant force pushing firms to outsource.

Unions and size-based regulations. Third, the upward-sloping labor supply curve could also be the result of union wage-setting or size-based labor market regulations. This concern is particularly relevant in France where firms with more than 50 employees face a number of legal obligations that may increase the cost of labor, including accepting a union delegate.²⁶ Our theory applies equally well if the labor supply curve is partly generated by union wage-setting or size-based regulations. Yet, we ask whether these institutions in France are likely to increase the cost of labor so much as to lead firms to outsource.

We use a regression discontinuity design to test the impact of unions and size-based regulations on outsourcing. Figure 20(a) in Appendix B.4 shows that French firms indeed bunch at the 50 employees threshold. However, Figure 20(b) in Appendix B.4 shows that there is no statistically significant nor economically meaningful discontinuity in the outsourcing share around the 50 employees threshold. If anything, consistently with Figure 5(b), the outsourcing-size gradient seems to flatten above 50 employees. We conclude that union wage-setting or size-based regulations do not impact outsourcing behavior significantly in France.²⁷

Having proposed reduced-form evidence supporting the key predictions of our theory, we turn to our general equilibrium quantitative exercises.

²⁶They must (i) form a committee that represents the interests of the employees to the management of the firm (“*comité d’entreprise*”), (ii) form a committee that monitors health and safety at work (“*comité d’hygiène, de sécurité et des conditions de travail*”), (iii) ratify an agreement that specifies what share of profits employees receive (“*accord de participation*”), (iv) maintain a monthly record of all hires and separations with the French administration (“*déclaration des mouvements de main-d’oeuvre*”), (v) establish a workforce-saving plan if they lay off more than 10 employees within a month (“*plan de sauvegarde de l’emploi*”), (vi) accept a union delegate, in which case annual wage bargaining takes place (“*délégué syndical*”), (vii) establish a plan to hire late-career employees (“*plan de génération et plan senior*”).

²⁷This result contrasts with Bertrand et al. (2020) who find substantial effects of firing restrictions on outsourcing in India. In addition to differences in the overall economic and institutional environment, the practical magnitude of firing costs imposed by size-based regulations may differ between both countries.

4 Extended model and estimation

To structurally evaluate the impact of outsourcing on the economy, we first enrich our environment along several dimensions. Next, we discuss the estimation strategy.

4.1 Quantitative setup

To capture the idea that goods producers may expand their human resource departments to hire more in-house, we let goods producers exert any recruiting effort—or post any number of vacancies— v in each market s . This recruiting effort comes at a convex cost $c_s(v) = c_{0s}v^{1+\gamma}$ for $\gamma > 0$. When $\gamma \rightarrow +\infty$ we recover the model of Section 2. The number of workers a goods producer attracts and retains now reflects its vacancy share:

$$n_s(w, v) = \frac{(1 + k_s)e_s}{(1 + k_s(1 - F_s(w)))^2} \cdot \frac{v}{V_s}, \quad (16)$$

where the equilibrium number of vacancies in market s satisfies

$$V_s = \mathcal{V}_s + M^G \sum_{\mathbf{o}: o_s=0} \int v_s(z, \boldsymbol{\varepsilon}, \mathbf{o}) \Omega(\mathbf{o}|z, \boldsymbol{\varepsilon}) \Gamma(dz, d\boldsymbol{\varepsilon}). \quad (17)$$

\mathcal{V}_s is the measure of vacancies posted by contractor firms. $v_s(z, \boldsymbol{\varepsilon}, \mathbf{o})$ denotes the number of vacancies posted by a firm with productivity z , outsourcing costs $\boldsymbol{\varepsilon}$ and decision $\mathbf{o} = \{o_s\}_s \in \{0, 1\}^S$. $\Omega(\mathbf{o}|z, \boldsymbol{\varepsilon})$ denotes the share of firms with productivity z and outsourcing costs $\boldsymbol{\varepsilon}$ that choose outsourcing bundle \mathbf{o} . $\Gamma(z, \boldsymbol{\varepsilon})$ is the joint cumulative distribution function of $(z, \boldsymbol{\varepsilon} \equiv \{\varepsilon_s\}_s)$.

In equilibrium, the number of workers that goods producer z attracts and retains is no longer a function of $\Gamma(z, \boldsymbol{\varepsilon})$ alone because goods producers with different outsourcing bundles \mathbf{o} post different wages at the same productivity z . Instead, this number of workers is a function of $\Upsilon_s(z, \boldsymbol{\varepsilon}, \mathbf{o}) = F_s(w_s(z, \boldsymbol{\varepsilon}, \mathbf{o}))$, where the wage offer distribution F_s is given in equation (49), Online Appendix G.2.

We impose a Cobb-Douglas revenue function nested in a decreasing returns upper tier:

$$R(z, \{n_s\}_s) = \left(z \prod_{s=1}^S n_s^{a_s} \right)^\rho, \quad \sum_{s=1}^S a_s = 1. \quad (18)$$

To relate our results to labor shares as simply as possible, we interpret decreasing returns to scale as arising from an underlying constant returns, Cobb-Douglas production function that uses capital and labor of different skills: $R(z, \{n_s\}_s) = k^{1-\rho} \left(z \prod_{s=1}^S n_s^{a_s} \right)^\rho$. For simplicity, we assume that goods producers require exactly one unit of capital to operate, leading to the revenue function in equation (18) when $k = 1$. We denote by η the cost of this unit of capital.

Goods producers then solve:

$$\begin{aligned} \pi(z, \{\varepsilon_s\}_s) &= \max_{\substack{\{n_s\}_s, \{v_s\}_s, \\ \{w_s\}_s, \{o_s\}_s}} R(z, \{n_s\}_s) - \sum_{s=1}^S \left\{ [(1 - o_s)w_s + o_s p_s \varepsilon_s] n_s + (1 - o_s) c_s(v_s) \right\} - \eta \\ \text{s.t. } n_s &= n_s(w_s, v_s) \text{ as per (16) if } o_s = 0. \end{aligned} \quad (19)$$

We specify a Cobb-Douglas matching function:

$$\mathcal{M}_s = \mu_s(m_s(u_s + \zeta_s(1 - u_s)))^\xi V_s^{1-\xi}. \quad (20)$$

ζ_s is the relative search intensity of employed workers: $\lambda_s^E = \zeta_s \lambda_s^U = \frac{\zeta_s \mathcal{M}_s}{m_s(u_s + \zeta_s(1 - u_s))}$. μ_s is the matching efficiency in market s .

The joint distribution of productivity and outsourcing costs (z, ε) is lognormal with respective standard deviations ν, σ and correlation ι . We normalize the log means to zero as they are not separately identified from $\{\tau_s\}_s$ and $\{b_s\}_s$. We interpret ε_s as an iceberg trade cost.

To capture the idea that contractors may specialize in screening and recruiting activities, we endow them with a possible comparative advantage in hiring. Their recruiting cost function is $c^C(v) = c_{0s}(\bar{c}^C)^\gamma v^{1+\gamma}$. The relative marginal cost $\bar{c}^C \leq 1$ lets the model match the outsourcing wage penalty together with size differences between contractors and goods producers by shifting the labor supply curve of contractors.

We solve the model for three skill types $S = 3$. To focus on low skill outsourcing, we impose that high skill workers ($s = 3$) are never outsourced: $\tau_3 = 0$. We also impose that only one type of low skill workers can be outsourced, $\tau_1 > 0, \tau_2 = 0$. We interpret the low skill workers who cannot be outsourced as “core” workers ($s = 2$), and those that can as “service” workers ($s = 1$). Core and service low skill workers share the same fundamentals δ_s, ζ_s . We relegate additional derivations to Online Appendix G.1 and computation details to Online Appendix G.4.

4.2 Estimation strategy

We set a quarterly frequency. We define low and high skill worker groups as revealed by the occupations of workers.²⁸ Non-employment is our primary measure of ‘unemployment’ in the model to capture steady-state flows into employment from individuals reported out of the labor force.

We estimate the model in three steps. We estimate a first group of parameters that can be directly mapped to data. Second, we set one parameter to an external value. In the third step, we estimate the remaining parameters jointly with a MSM estimator. To be consistent with our use of the full sample between 1997 and 2007 to construct most moments, we interpret our baseline parameter estimates as reflective of the French economy at the 2002 midpoint.

First, we identify the parameters $\{\delta_s, \zeta_s\}_{s=1}^3$ from labor market flows. The time-aggregated employment-to-non-employment transition rate EN_s is equal to the job losing rate parameter δ_s . The non-employment-to-employment transition rate NE_s is equal to the endogenous offer rate λ_s^U from non-employment. Online Appendix G.6 relates the employment-to-employment transition rate EE_s to the arrival rate $\lambda_s^E = k_s \delta_s$:

$$\frac{\text{EE}_s}{\text{EN}_s} = \frac{(1 + k_s) \log(1 + k_s) - k_s}{k_s}.$$

We recover $\zeta_s = \lambda_s^E / \lambda_s^U$.

²⁸We rank 2-digit occupations by their average wage. We compute the mean predicted wage of a worker by interacting the mean occupational wages and the time spent by the worker in each occupation. A worker is low skill if their predicted wage is below median, and high skill if above.

Second, we set the elasticity of the matching function ξ to a common value of $\xi = 0.5$ since our data does not let us estimate it credibly (Petrongolo and Pissarides, 2001).

Third, we jointly estimate the remaining parameters $\{\mu_s\}_{s=1}^3, \{b_s\}_{s=1}^3, \{a_s\}_{s=2}^3, \rho, \gamma, \tau_1, \nu, \sigma, \iota, m_1, M^G, \eta, \bar{c}^C$ by MSM. While the parameters are jointly identified, we provide an intuitive argument for identification. We confirm our argument numerically in Figure 21 in Appendix C.

Equation (20) reveals that the matching function efficiency for skill s , μ_s , has a direct impact on the non-employment-to-employment transition rates NE_s which we target for each skill. By shifting unemployment benefits conditional on wages, the parameters b_s affects the replacement rate for each skill type which we target. We inform the relative demand for high skill workers a_3 using the skill premium. We identify the relative demand a_2 for service vis-à-vis core low skill workers imposing that they earn identical average wages. The curvature in the revenue function ρ shifts the aggregate labor share.

We identify the measure of goods producers M^G with their average employment, and the measure of contractors M^C with their employment relative to goods producers. The capital fixed cost η affects the fraction of firms with low employment. The parameters γ and ν jointly determine the dispersion in size and value added. When there is less curvature in recruiting costs, productive firms can hire more and dispersion in firm size increases. Conditional on size, the dispersion in productivity raises dispersion in value added. Thus, we target the standard deviation of log firm size to inform γ , and the standard deviation of log value added to inform ν .

The outsourcing cost τ_1 determines how much outsourcing there is in the economy by directly affecting the returns to outsourcing. Thus, we target aggregate outsourcing expenditures as a fraction of wages. We determine the measure of service low skill service workers m_1 by matching the employment outsourcing share for service workers. The dispersion in outsourcing costs σ determines the fraction of goods producers who have an outsourcing share below 10%. To inform the correlation between productivity and outsourcing costs ι , we target the OLS relationship between the outsourcing share and value added. Importantly, we use a cross-sectional moment rather than a within-firm moment. We use the recruiting comparative advantage of contractors \bar{c}^C to match the outsourcing wage penalty jointly with contractor firm size. We provide additional details in Online Appendix G.5.

4.3 Estimation results and identification

Table 11 in Appendix C summarizes our estimation results. Our estimates fall within conventional ranges found in the literature. The revenue function curvature parameter is $\rho = 0.92$. Our estimate of a_3 implies that the marginal product of labor for high skill workers is two and a half times as large as for low skill workers at equal employment shares. The curvature in the vacancy cost $\gamma = 10.1$ is towards the high end of values reported in the literature. Low employment dispersion relative to value added in the data together with ρ being close to 1 requires a high degree of recruiting cost curvature. We find that 37% of workers in the French economy are exposed to outsourcing.

Our estimate of τ_1 implies that outsourced workers are 35% less productive than in-house workers. By jointly targeting outsourcing expenditure and employment shares, the estimation requires that τ_1 and the outsourcing price p_1 move in opposite directions. Increasing τ_1 towards 1 implies that the

outsourcing price p_1 falls. Demand rises more than one-for-one because of decreasing returns to scale. In general equilibrium, this excess demand is incompatible with the outsourcing employment share. If workers at contractors were as efficient as in-house employees, we should observe a higher outsourcing employment share relative to the expenditure share. Thus, we infer that the data favor the cost-saving view relative to the comparative advantage view.

We find that contractors, however, are more efficient than goods producers at recruiting workers. The marginal recruiting cost of contractors is 55% of the marginal recruiting cost of goods producers. This difference in marginal recruiting costs is necessary to rationalize that contractors are larger than goods producers despite paying lower wages.²⁹ Thus, the data favors a view in which contractors specialize in screening and recruiting activities rather than production activities.

How well are parameters identified? To answer this question, Figure 21 in Appendix C plots both the simulated moment (univariate identification) and the loss function (multivariate identification) as we vary the parameter close to its estimated value. Most of the parameters are well identified locally. Moment deviations are steep functions of parameter deviations, and the loss function is peaked around 0.³⁰ Overall, Figure 21 confirms our identification argument and supports our estimation strategy.

4.4 Over-identification

We propose three over-identification checks that relate to the three main predictions of our theory. The first relates to selection into outsourcing. The second relates to the productivity effect. The third relates to the distributional effect.

We start by verifying whether the estimated model accounts for selection into outsourcing. We have targeted the cross-sectional OLS coefficient in column (2) from Table 6, Appendix B.1 to inform σ . However, focusing on within-firm changes (column 3) as well as instrumenting for firm revenue productivity (column 6) affects this coefficient. While these are non-targeted moments, can the estimated model rationalize these differences?

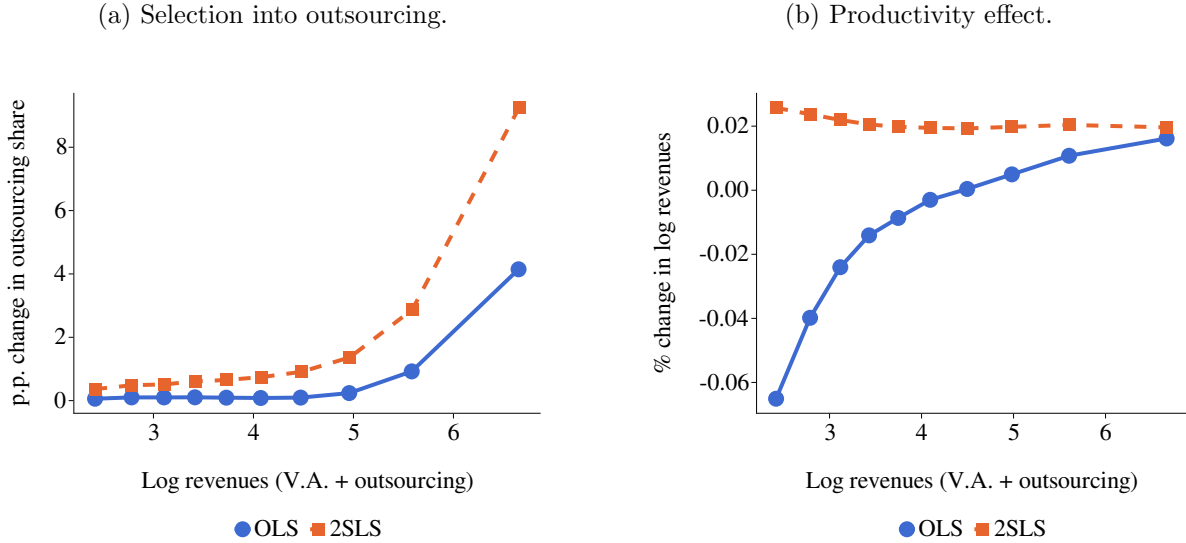
The solid blue line in Figure 8(a) displays the model equivalent of the coefficient in column (3), from the following experiment. Consider a one standard deviation shock to revenue productivity z . We estimate that z and ε are positively correlated—more productive firms face larger outsourcing costs. This positive correlation implies that the increase in z is also associated with an average increase in ε . For every firm, we compute the change in the outsourcing share following the joint change in (z, ε) and project it on the associated change in value added. We then display the resulting OLS coefficient in the model by decile of initial value added. We cannot aggregate these OLS effects in the model because we cannot reliably estimate the incidence of shocks firm by firm. Nevertheless, our empirical estimate of 1.82 lies in the middle of the OLS effects in the model that range from 0.06 to 4.15.

In the model, the within-firm OLS coefficient conflates the change in revenue productivity z with the associated change in outsourcing costs ε . We mimic the instrumental variable strategy from Table 6, column (6) as follows. We interpret the export demand instrument as removing the increase in ε associated with the increase in z . We then only shift z to compute the change in the outsourcing

²⁹The need for such a shift in the labor supply curve of contractors to rationalize the size gap is not specific to our environment. It would be necessary in any model with a firm-specific labor supply curve to generate wage premia.

³⁰Lest we simulate the model on a nineteen-dimensional hypercube, we cannot however guarantee global identification.

Figure 8: Selection into outsourcing and the productivity effect in the estimated model.



Note: Panel (a): OLS (solid blue) and 2SLS (dashed orange) coefficients in the estimated model. OLS coefficient computed by projecting the change in the outsourcing share on the change in log value added, following a one standard deviation Δz increase in z and the corresponding change in $\Delta \varepsilon = \frac{\sigma}{\nu} \Delta z$. 2SLS coefficient computed by only increasing z by one standard deviation. Panel (b): OLS (solid blue) and 2SLS (dashed orange) coefficients in the estimated model. OLS coefficient computed by projecting the change in log value added on the change in the outsourcing share, following a one standard deviation $\Delta \varepsilon$ decrease in ε and the corresponding change in $\Delta z = \frac{\nu}{\sigma} \Delta \varepsilon$. 2SLS coefficient computed by only increasing ε by one standard deviation.

share and value added instead of the joint shift in (z, ε) . The dashed orange line displays our results. Consistently with the data, the model counterpart of the 2SLS estimate is much larger than the OLS estimate. This ordering occurs because of the positive correlation between z and ε . When revenue productivity z rises alone, firms are more inclined to increase outsourcing than when the cost of outsourcing increases simultaneously. Quantitatively, the magnitude of the 2SLS coefficient lies between 0.36 and 9.25 depending on the incidence of shocks in the model, consistently with our point estimate of 3.41.

Our second over-identification exercise asks whether the estimated model accounts for the untargeted productivity effect. We mimic the OLS and 2SLS coefficients from Table 7 similarly to selection into outsourcing. We consider a negative one standard deviation shock to the idiosyncratic outsourcing cost ε_s of the firm. We then decrease revenue productivity z accordingly for the OLS coefficient, or leave it unchanged for the 2SLS coefficient.

Figure 8(b) displays the within-firm model counterparts of the OLS and 2SLS coefficients from Table 7, columns (1) and (3), by initial decile of firm value added. In the model, the OLS coefficient ranges from -0.07 to 0.01 depending on the incidence of shocks, consistently with our point estimate that is near zero. The model struggles to generate a high enough 2SLS productivity effect, which remains between 0.02 and 0.03. By contrast, our empirical estimate is close to 0.08. This limitation implies that the model will likely under-predict the employment response from outsourcing in the aggregate.

Taking stock, the estimated model is quantitatively consistent with three out of four of these non-targeted moments, and qualitatively consistent with the fourth. Hence, we conclude that the estimated model provides an empirically plausible account of selection into outsourcing and the productivity

Table 1: Wage dispersion and labor market statistics in the model and the data.

	Data			Model		
	A. Aggregates					
St. dev. firm premia						
AKM firm premia	0.14			0.17		
Within-skill wage st. dev.				0.16		
	B. By producer type					
	Contractors	Goods prod.	Diff.	Contractors	Goods prod.	Diff.
Firm premia	-0.14	0.00	-0.14	-0.13	0.00	-0.13
Frac. hires from emp.	0.46	0.52	-0.06	0.44	0.56	-0.12
Churn	0.11	0.03	0.08	0.10	0.08	0.02
Net poaching	-0.02	0.00	-0.02	-0.01	0.00	-0.02

effect.

Our third over-identification exercise verifies whether the estimated model accounts for the distributional effects of outsourcing through the job ladder in the model. The estimation targets the outsourcing wage penalty, but does not target wage dispersion. Instead, wage dispersion emerges endogenously from productivity dispersion and worker flows, that are both disciplined with data. Wage dispersion is linked to other non-targeted labor market statistics such as the employment hire share or net poaching through the job ladder structure of the model.

Table 1 indicates that the model matches these non-targeted moments well. To estimate firm wage premia consistently with our theory, we run equation (12) similarly to Section 3.3. Using the AKM decomposition (12), we estimate the within-skill standard deviation of firm wage premia to be 0.14 in the data. When we run the same AKM decomposition in our estimated model, the within-skill standard deviation of log wage premia is 0.17. It is close to the within-skill standard deviation of wages (0.16) because workers only differ by skill in the model.

In the data, contractors hire 46% of their workers from employment, while goods producers hire 52% of their workers from employment. These hiring patterns are closely mirrored in the model: contractors hire 45% of their workers from employment, while goods producers hire 56% of their workers from employment. Contractors have relative net poaching of -2% in the model and in the data. Contractors also have more churn in the model, qualitatively consistent with the data. We conclude that our job ladder structure provides a consistent account of wage dispersion and systematic labor market differences between contractors and goods producers. Having validated the core structure of the model, we next estimate the shocks to outsourcing that drive our counterfactuals.

4.5 Outsourcing demand and supply shocks

Our counterfactual exercises assess the impact of outsourcing on wages and employment. As in any supply and demand framework, outsourcing demand and supply shocks are expected to have different effects on wages and employment. Therefore, we consider both types of shocks. Supply shocks

Table 2: Industry-level correlations between changes in outsourcing share and changes in wages, employment, value added, profits and the labor share.

	(1)	(2)	(3)	(4)	(5)	(6)
	Wage	Employment	Value added	Contractor size	Profits	Labor share
Out. share (data)	-0.02	0.06	0.04	-0.16	0.06	0.00
Out. share (model)	-0.02	0.06	0.04	-0.15	0.05	0.00
Obs.	1150	1150	1150	1150	1150	1150

Linear regressions of changes in log mean wage, log employment, log value added, log contractor size, log residual profits and labor share on changes in the outsourcing share at the commuting zone level. Changes defined between 1996, 2001 and 2007. Residual profits defined as value added net of wage bill. Coefficients normalized to represent the effect of a 5 p.p. increase in the outsourcing share. Includes year dummies.

are changes in the productive and recruiting comparative advantage of contractors τ_1, \bar{c}^C , and the measure of contractors M^C . Demand shocks are changes in the correlation between productivity and outsourcing costs ι .

We disentangle demand and supply shocks by matching cross-industry regressions of wages and employment, value added and contractor size on the outsourcing share. We report our results in columns (1-4) of Table 2. We use labor shares and accounting profits in columns (5-6) as over-identification checks. We also report identification plots for each shock in Figure 22, Appendix C. We discipline the overall magnitude of our shocks by replicating the increase in the aggregate expenditure share on outsourcing from Figure 2 between 1997 and 2007. Since our baseline parameter estimates represent the 2002 French economy, we allow τ_1, \bar{c}^C, M^C and ι to either increase or decrease in fixed proportions to match the aggregate expenditure share in 1997 and 2007.

We find evidence of demand shocks as industries with stronger increases in the outsourcing share experience a rise in value added. We report our estimates in Table 12, Appendix C. The 0.04 correlation between industry-level value added changes and outsourcing changes implies that the correlation between productivity and outsourcing costs drops from 0.17 to 0.11 between 1997 and 2007, driving up demand for outsourcing.

We also infer strong supply shocks. The 0.06 cross-sectional positive correlation with employment implies that contractors experience strong improvements in their recruiting comparative advantage. \bar{c}^C drops from 1.03 in 1997 to 0.31 in 2007. Our baseline estimate for 2002, 0.55, lies in the middle of these values. The 16% decline in contractor size requires that more than twice as many contractors operate in 2007 relative to 1997 given the improvement in recruiting technology. To fully match the 0.02 decline in wages, we also infer a relative decline in the productive comparative advantage of contractors, which falls from 0.78 to 0.54.³¹ Columns (5-6) in Table 2 show that the model also matches non-targeted industry-level changes. It matches the weak correlation between outsourcing and the labor share as well as the strong correlation between accounting profits and outsourcing.

³¹Our model is isomorphic to one in which there is free entry of contractors, and the entry cost varies over time. Thus, the rise in the measure of contractors is consistent with a decline in contractor entry costs together with an improvement in their recruiting comparative advantage. In a broader framework in which contractors would be heterogeneous in τ_1 , one could also rationalize the decline in τ_1 through adverse selection following the decline in entry costs. For simplicity, we instead directly estimate the decline in productive comparative advantage τ_1 required to match the data.

5 The impact of outsourcing on inequality and output

We are now ready to investigate the impact of the rise in outsourcing on inequality and aggregate output. We compare steady-states of the estimated model, and interpret our results as the effect of outsourcing on the French labor market in the decades between 1997 and 2016. We index our counterfactuals by the aggregate outsourcing expenditure share.

As emphasized in Section 3.2, our firm-level expenditure data stops in 2008. We use the model to ask what is the effect of outsourcing on the French labor market by 2016. We choose the midpoint between the 2007 outsourcing expenditure share and our extrapolation in Figure 3(a) to remain conservative. To represent the 2016 economy in our model, we use demand and supply shocks in the same proportions as between 1997 and 2007, and their overall magnitude is set to match the rise in the outsourcing expenditure share. In our figures, we represent the 2007 counterfactual economy with a vertical dotted line.

5.1 The impact of outsourcing on service workers

We start with the impact of outsourcing on wages and earnings of service workers, displayed in Figure 9. The fraction of service workers employed at contractor firms rises by 21 p.p. by 2007, as shown in panel (a). This increase closely tracks the one observed in the data in Figure 3(a) although it is non-targeted. Thus, the link between outsourcing expenditures and outsourcing employment in the aggregate coincides in the model and in the data. By 2016, the fraction of low skill service employment at contractors has risen by 37 p.p.

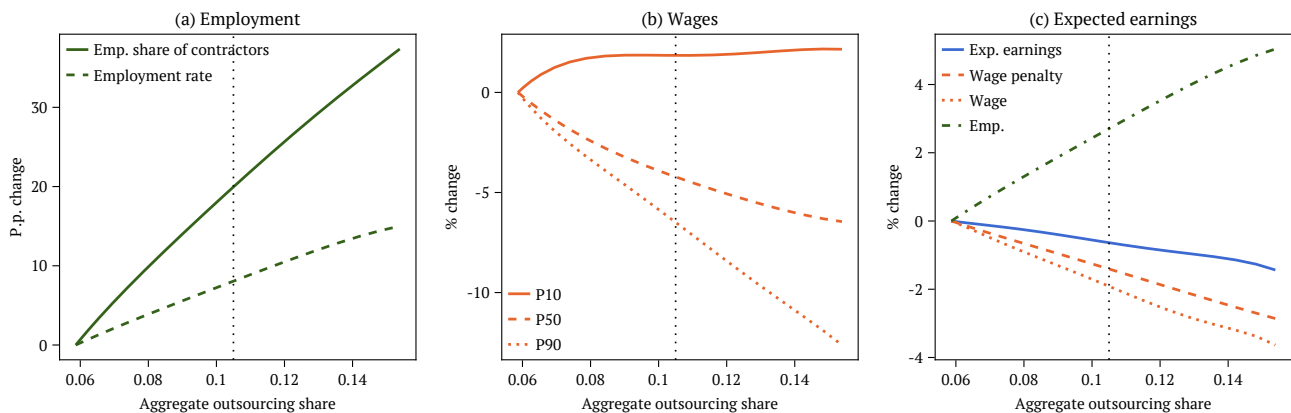
The reallocation of low skill services workers towards contractors precludes to adverse impacts on earnings because of the outsourcing wage penalty. Yet, low skill service workers also experience strong employment gains because outsourcing lowers the effective cost of labor in the aggregate. Panel (a) reveals that outsourcing lowers their non-employment rate by 14.9 p.p., down from 25 p.p in 1997.

Wages of low skill service workers respond ambiguously in general equilibrium as shown in panel (b). Rising outsourcing implies tougher competition for workers at the bottom of the job ladder where contractors operate. Therefore, wages in the bottom decile rise by 2.1%. At the same time, outsourcing removes the highest paying jobs from the labor market. In addition, the reallocation towards low-paying contractor firms drags down the entire wage distribution. Thus, median wages drop by 6.6%, and wages in the top decile fall by 12.6%.

Expected earnings of low skill service workers thus reflect several possibly offsetting forces. We present an exact decomposition of expected earnings in Appendix A.17. The reallocation towards contractor firms lowers earnings because of the outsourcing wage penalty. The general equilibrium response of employment raises earnings by pulling more workers out of non-employment. General equilibrium changes in the wage distribution of employed workers are mostly detrimental to earnings, except at the bottom of the job ladder.

Quantitatively, panel (c) shows that the reallocation of low skill service workers towards contractor firms results in 2.8% earnings losses in expected earnings by 2016 once combined with the 14% outsourcing wage penalty. The general equilibrium response of the wage distribution in panel (b) aggregates to a further 3.8% decline in earnings. The rise in low skill service employment is the strongest

Figure 9: The effects of outsourcing on wages, earnings and welfare of service workers.



Note: Demand and supply shocks calibrated to match the rise in outsourcing expenditure share in France from 1997 to 2007. All outcomes are shown as a function of aggregate outsourcing expenditures relative to the aggregate wage bill on the x-axis. Leftmost point corresponds to 1997 economy. Vertical dotted line corresponds to counterfactual 2007 economy. Rightmost point corresponds to counterfactual 2016 economy. Panel (c): expected earnings and decomposition as in Appendix A.17.

general equilibrium feedback. As outsourcing rises, effective labor market frictions decline and so does the cost of labor. Employment of low skill service workers rises by 13%. Panel (c) shows that it improves expected earnings by 5% by 2016.

Once we aggregate these three channels, we find that low skill service workers ultimately experience a net 1.5% decline in expected earnings. Employment gains are dominated by the partial and general equilibrium declines in wages. We show in Appendix A.17 that expected earnings coincide with welfare given our assumption of a small discount rate. Thus, low skill service workers lose from outsourcing.

The decomposition in panel (c) highlights that general equilibrium effects are critical to evaluate the distributional impact of outsourcing on the economy. The partial equilibrium impact alone would overstate welfare losses of low skill service workers and miss important general equilibrium margins of adjustment. Our model lets us quantify these margins and aggregate them into welfare.

5.2 The impact of outsourcing on sorting and wage inequality

Are welfare losses for service workers accompanied by sizeable changes in wage inequality? A rapidly expanding literature on inequality and labor market sorting hypothesizes that the outsourcing-induced reallocation of low skill workers to low paying firms as in Section 5.1 increases labor market sorting and inequality (see for instance Goldschmidt and Schmieder (2017) and Song et al., 2018).

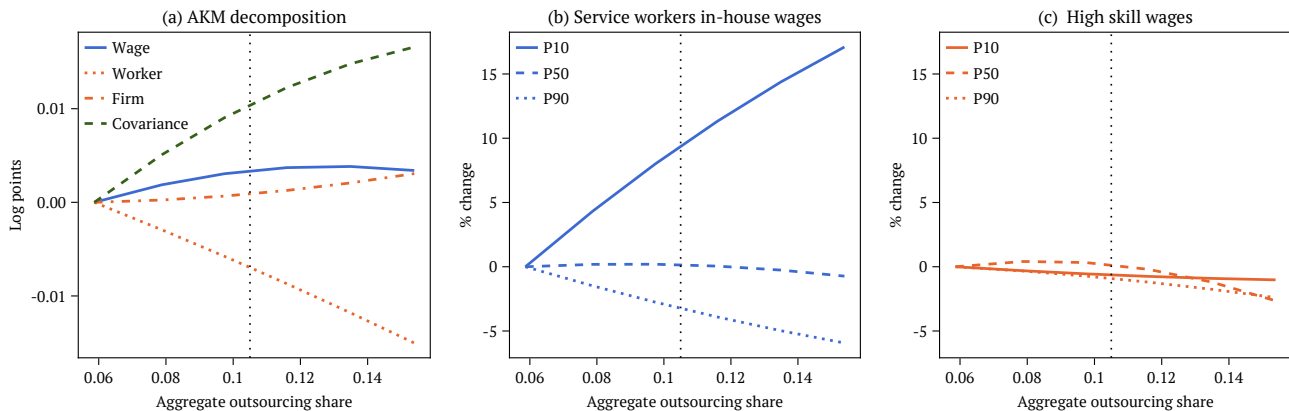
We evaluate this hypothesis with our structural model. To stay as close as possible to previous work, we run an AKM regression in our estimated model similarly to Section 3.3:

$$\log w_s(z) = \alpha_s + \varphi(z) + \varepsilon_s(z). \quad (21)$$

Our main measure of labor market sorting is the correlation between worker and firm effects, $\text{Corr}[\alpha_s, \varphi(z)]$. The conditional random mobility assumption necessary for identification is satisfied in our environment, although the wage function is not exactly log-linear. Nevertheless, we obtain a close fit between the estimated firm effects and actual wages in the model, as shown in Figure 23, Appendix D.

Perhaps surprisingly, we find that outsourcing has a limited impact on wage inequality in general

Figure 10: The impact of outsourcing on labor market sorting and wage inequality.



Note: Decomposition of log wages into worker and firm effects following equation (21) in the model. Panel (a): variance decomposition of log wages following equation (22). Panel (b): low skill service wages at goods producers. Panel (c) high skill wages.

equilibrium. Panel (a) in Figure 10 shows that the variance of log wages barely rises between our model 1997 and 2016 economies. In light of the conventional wisdom that outsourcing should increase wage inequality, we use the AKM decomposition in equation (21) to perform an exact employment-weighted variance decomposition and understand the limited impact of outsourcing on inequality:

$$\text{Var}[\log w_s(z)] = \text{Var}[\alpha_s] + \text{Var}[\log \varphi(z)] + 2\text{Cov}[\alpha_s, \varphi(z)] + \text{Var}[\varepsilon_s(z)]. \quad (22)$$

We find that outsourcing strongly increases labor market sorting. Panel (a) reveals that the covariance between the worker component of wages, α_s , and the firm component of wages, $\varphi(z)$, increases by nearly 0.015 between 1997 and 2016. This rise is substantial given that we only have three worker types. Our results are consistent with the hypotheses in Goldschmidt and Schmieider (2017) and Song et al. (2018). The correlation between worker and firm effects increases by 23.3 p.p. by 2016, indicative of a substantial rise in labor market sorting.

However, our general equilibrium environment uncovers a force that offsets the rise in labor market sorting: the variance of worker effects declines so much with outsourcing that it offsets the rise in the covariance between AKM worker and firm effects. To understand this decline, three observations are necessary.

The first observation is that the AKM decomposition (21) identifies worker effects α_s by comparing wages of different workers employed at the same firm. Thus, low skill service employees at goods producers identify worker effects because they are the only ones who can be compared to other workers within the same firm. Only low skill service workers are employed at contractors. To understand the decline in the variance of worker effects, we thus need to understand how wages of low skill service and high skill workers change at goods producers.

The second observation is that outsourcing leads to pro-competitive effects at the bottom of the job ladder. Wages of low skill service workers employed at goods producers respond more strongly than wages of all low skill service workers. Panel (b) of Figure 10 indicates that wages at low-paying goods producers rise by over 15%. Median wages remain constant, while wages at the top decline by 5%. Compared to wage changes at all firms in Figure 9(b), these changes are more favorable to service workers because they net out reallocation towards contractors.

The third observation is that wages of high skill workers remain largely stable as outsourcing rises. Panel (c) of Figure 10 shows that wages in the bottom decile, at the median and in the top decile for high skill workers do not change as outsourcing rises. That wages of high skill workers remain stable despite the complementarity between skills in production and the expansion of low skill service employment at large productive firms (the productivity effect) is related to rent-sharing. We return to this point in the next section.

These three observations imply that the AKM decomposition (21) infers that low skill service workers are better paid compared to their high skill coworkers in 2016 relative to 1997.³² Thus, the variance of worker effects declines, and wage inequality remains stable despite a substantial rise in labor market sorting.

Our results thus caution against inferring that a rise in labor market sorting necessarily implies a rise in wage inequality. Our environment highlights that outsourcing causes both rising sorting and pro-competitive effects. Wage inequality remains largely constant as both forces offset each other.

5.3 The impact of outsourcing on rent-sharing

Labor market sorting is only one aspect of inequality. Sorting captures differences in labor market opportunities between different workers. Our environment also captures how the marginal product of labor is split between workers and firms, which we define as rent-sharing. We now investigate how outsourcing changes rent-sharing at goods producers who hire in-house.

Our main measure of rent-sharing is the markdown paid by goods producers. The markdown is defined as the ratio of wage to the marginal product of labor: $\text{markdown}_s(z) = w_s(z)/\text{MPL}_s(z)$. A markdown close to 0 implies that workers recover a small fraction of the rents from production. A markdown close to 1 implies that workers obtain nearly all the rents from production.

A related measure of rent-sharing is the firm-level labor share:

$$\text{LS}(z) = \sum_{\text{in-house } s} \frac{n_s^*(z)w_s(z)}{\text{VA}(z)} = \sum_{\text{in-house } s} \text{markdown}_s(z) \frac{n_s^*(z)\text{MPL}_s(z)}{\text{VA}(z)}. \quad (23)$$


Simple accounting ensures that the labor share is a weighted average of markdowns. These weights need not add up to unity in presence of decreasing returns to scale. Markdowns are the conceptually exact metric of rent-sharing, but they are difficult to estimate at the firm-by-skill level. The firm-level labor share is a coarser measure of rent-sharing but is easier to measure.

Our micro-foundation for labor market power delivers dispersion in markdowns and labor shares across firms. We cannot measure markdowns directly, but we can contrast our model with the data for labor shares. Panel (a) in Figure 11 displays the distribution of labor shares across firms. Despite being entirely non-targeted, the distribution of labor shares in the model is close to the data although not entirely right-skewed enough.

Our environment also captures the co-movement of labor shares with wages and value added. Panel (b) shows that labor shares fall with firm wages, with a close fit between model and data except perhaps at the very top. Similarly, panel (c) indicates that labor shares decline with firm value added in both

³²This implication would continue to hold in a richer environment in which contractors also employ high skill workers, as long as they employ relatively less high skill workers than goods producers

Figure 11: Rent-sharing in the labor market.



Note: Panel (a): employment-weighted distribution of labor shares. Beta distribution fitted to both model and data for comparability. Panel (b): Labor share by average wage. Panel (c): Labor share by value added. Only goods producers. Results with all firms in Figure 24, Appendix D.

model and data. These results suggest that workers extract a smaller fraction of rents when working at highly productive large employers, as highlighted in Autor et al. (2020) and Gouin-Bonenfant (2022).³³

We unpack how different skill groups contribute to rent-sharing by investigating markdowns in our economy. Panel (a) in Figure 12 splits the average wage paid by goods producers to service workers into the marginal product of labor and the markdown according to the identity

$$\mathbb{E}[\log w_s(z)] = \mathbb{E}[\log \text{MPL}_s(z)] + \mathbb{E}[\log \text{markdown}_s(z)] \quad (24)$$

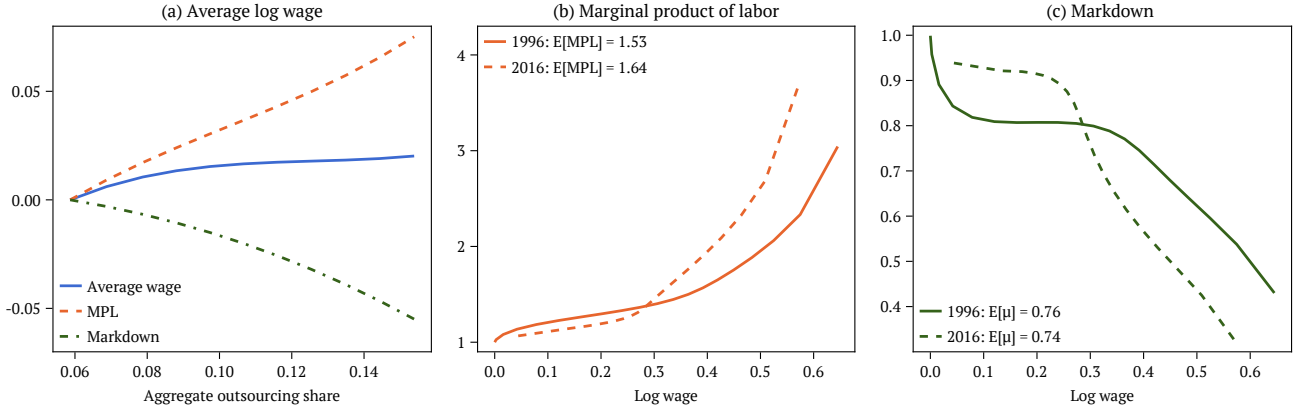
for service workers $s = 1$.

Among goods producers, the marginal product of labor rises with outsourcing. Panel (b) shows that as service workers reallocate towards contractors, goods producers employ fewer service workers and the marginal product of labor increases. This rise is particularly pronounced at the top of the job ladder that even fewer workers ever reach. Figure 25 in Appendix D confirms that the marginal product of labor only rises at goods producers who hire in-house. When we include goods producers who outsource, the marginal product of labor instead falls on average at highly productive firms, as firms switch from an upward-sloping to a vertical labor supply curve.

The rise in the marginal product of labor is offset by a strong equilibrium decline in the average markdown at goods producers. This decline results from the clockwise tilt in the markdown distribution depicted in panel (c). Markdowns at goods producers rise at low-paying firms. As contractors compete with goods producers for workers in the lower rungs of the job ladder, workers are able to extract a larger fraction of the rents from production. Given the option to cheaply outsource however, high-paying goods producers only find it profitable to hire in-house at low markdowns. This selection implies that markdowns fall at the higher end of the wage distribution. On net, the average markdown drops and workers do not recoup the increase in the marginal product of labor. Our results highlight that both reallocation and changes in rent-sharing are key to understand how outsourcing impacts wages in general equilibrium.

³³By contrast, models of wage dispersion based on compensating differentials do not generate within-market between-firm dispersion in labor shares or markdowns with standard revenue functions. See Online Appendix F.2 for details.

Figure 12: The impact of outsourcing on rent-sharing for service workers.



Note: Panel (a): decomposition of mean log wage for service workers at goods producers according to the identity (24). Panel (b): distribution of the marginal product of labor for service workers at goods producers who hire in-house. Panel (c): distribution of markdowns for service workers at goods producers who hire in-house. See Figure 26 in Appendix D for results with all firms.

The impact of outsourcing low skill service workers spills over to high skill workers through the supermodularity condition in Assumption (A). When employment of service workers rises at goods producers, so does the marginal product of labor for core low skill and high skill workers. Thus, our result that wages of high skill workers are largely stable as outsourcing rises (Figure 10) is surprising at first sight.

A more detailed investigation reveals that the marginal product of labor and markdowns also move in opposite directions for high skill workers.³⁴ Panel (a) in Figure 13 indicates that the average marginal product of labor rises moderately with outsourcing for high skill workers. Yet, the average markdown simultaneously drops, implying slightly declining wages for high skill workers.

The marginal product of labor indeed rises only at the top of the job ladder as shown in panel (b), where firms are outsourcing service workers. Yet, markdowns further decline because of the productivity effect. Outsourcing is akin to an increase in revenue productivity from the perspective of high skill workers. Thus, goods producers that outsource more move rightwards in the markdown distribution where markdowns are lower in equilibrium (Figure 11). Similarly to service workers, offsetting movements in markdowns and the marginal product of labor shape high skill wages in general equilibrium.

5.4 The impact of outsourcing on output and TFP

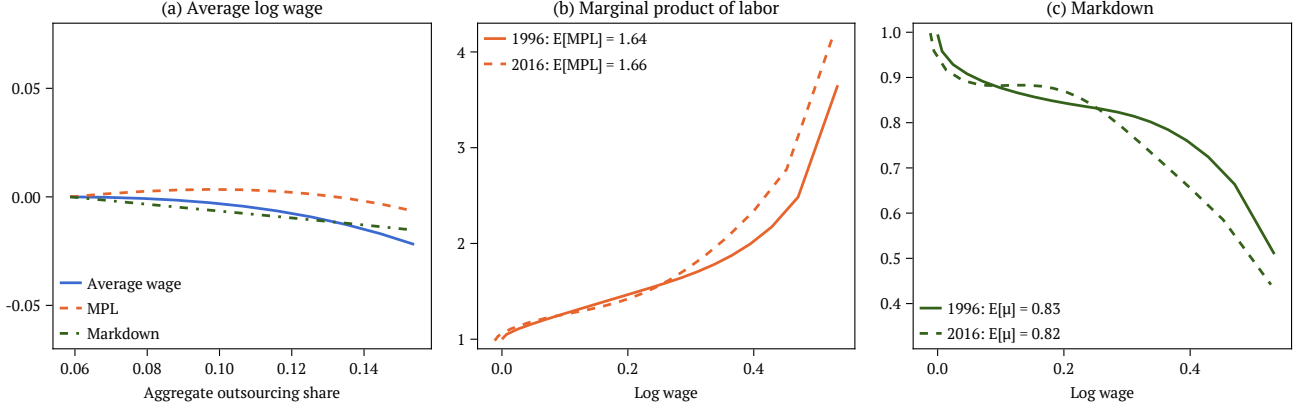
Despite the adverse effect of outsourcing on the welfare of low skill service workers and its subtle implications for sorting and rent-sharing, outsourcing may generate output and productivity gains. We describe the impact of outsourcing on aggregate output in Figure 14, panel (a).

We find that aggregate output rises by 6.1%, in line with strong productivity effects of outsourcing. This increase in output could in principle be driven by extensive margin adjustments—a rise in employment—or by intensive margin adjustments—a rise in TFP.

Output gains are entirely driven by the 10.3% increase in aggregate employment, in turn largely

³⁴Low skill core workers display similar patterns to high skill workers.

Figure 13: The impact of outsourcing on rent-sharing for high skill workers.



Note: Panel (a): decomposition of mean log wage for high skill workers at goods producers according to the identity (24). Panel (b): distribution of the marginal product of labor for high skill workers at goods producers. Panel (c): distribution of markdowns for high skill workers at goods producers.

due to low skill service workers. In fact, our results indicate that outsourcing has a negative 2% impact on aggregate TFP. Panel (b) reveals that this decline in turn masks two offsetting effects captured by an exact decomposition.

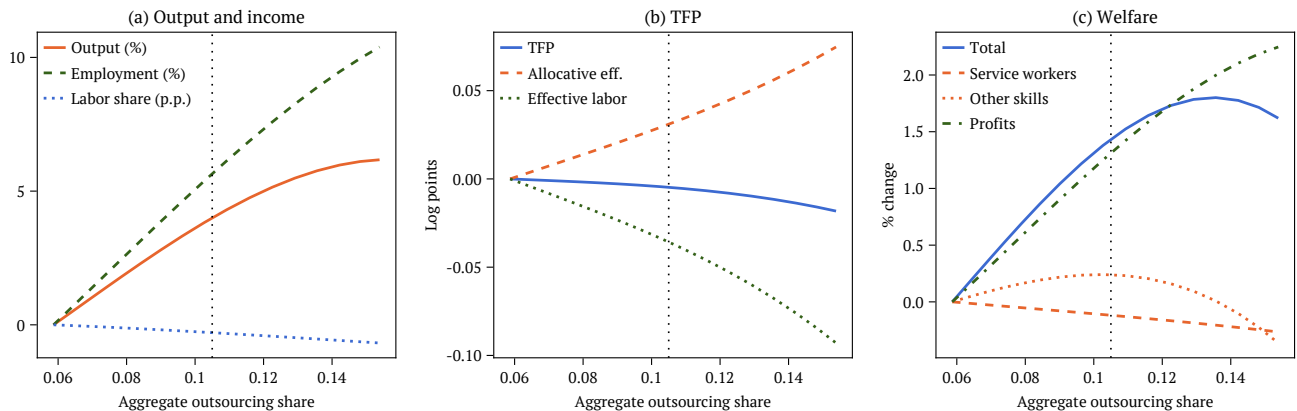
We define two labor aggregators consistent with our production function. The first one is simply the Cobb-Douglas labor aggregator that corresponds to the revenue function: $\hat{N} = \left(\prod_{s=1}^3 \bar{N}_s^{a_s} \right)^\rho$, where \bar{N}_s denotes total employment of skill s . The second aggregator is $\tilde{N} = \left(\bar{N}_1^G + \tau_1 \bar{N}_1^C \right)^{\rho a_1} \left(\prod_{s=2}^3 \bar{N}_s^{a_s} \right)^\rho$ where \bar{N}_1^G, \bar{N}_1^C denote aggregate employment of service workers at goods producers and contractors, respectively. \tilde{N} captures the idea that effective labor of low skill service workers is lower when more of them work at contractors if $\tau_1 < 1$. Changes in the ratio \tilde{N}/\hat{N} thus encode changes in effective labor. Our exact TFP decomposition reads:

$$\Delta \log \text{TFP} = \underbrace{\Delta \log \frac{M^G \mathbb{E}[R(z, \mathbf{n})]}{\tilde{N}}}_{\text{Allocative efficiency given effective labor}} + \underbrace{\Delta \log \frac{\tilde{N}}{\hat{N}}}_{\text{Effective labor}}. \quad (25)$$

Consistently with the idea that outsourcing reallocates labor towards highly productive firms that were constrained by labor market frictions, allocative efficiency rises by 7%. But the reallocation of low skill service workers towards less productive contractors drags TFP down by nearly 9%. Crucially, the comparative advantage view would predict a rise in effective labor. By contrast, the cost-saving view reveals that allocative TFP gains from outsourcing are muted by a strong reduction in effective labor.

Perhaps surprisingly, we find that outsourcing has little effect on the labor share. This result follows from two observations. First, high skill and core low skill workers experience modest wage and employment changes. Thus, they do not contribute strongly to changing the labor share through either changes in markdowns or weights as per the decomposition (23). Second, markdowns and weights move in opposing directions for low skill service workers. Markdowns decline (panel (a) in Figure 12). However, service workers are also reallocated towards contractors. Contractors locate

Figure 14: The impact of outsourcing on aggregate output and TFP.



Note: Demand and supply shocks calibrated to match the rise in outsourcing expenditure share in France from 1997 to 2007. All outcomes are shown as a function of aggregate outsourcing expenditures relative to the aggregate wage bill on the x-axis. Leftmost point corresponds to 1997 economy. Vertical dotted line corresponds to counterfactual 2007 economy. Rightmost point corresponds to counterfactual 2016 economy. Panel (b): TFP calculated as in equation (25).

at the bottom of the job ladder and thus have high labor shares. These two effects cancel out in the aggregate.³⁵ As a result, panel (c) reveals that all worker types experience weak or negative welfare changes, while profits rise and drive up aggregate welfare.

5.5 The impact of outsourcing with a binding minimum wage

Our analysis finds that low skill service workers ultimately lose from outsourcing. A natural question is whether standard labor market policy instruments can shield low skill service workers from the outsourcing wage penalty while maintaining most employment and output gains. We focus on the minimum wage as our main policy instrument. The minimum wage is a natural candidate since it maintains wages at any desired level. The counterpart is that the minimum wage may push up the price of labor so much that it deters vacancy creation and reduces employment, ultimately lowering output. We use our estimated model to explore whether outsourcing can benefit workers and firms equally with a simple minimum wage reform.

We conduct a minimum wage experiment, in conjunction with the same outsourcing demand and supply shocks as in Sections 5.1 to 5.4. The minimum wage is only binding for low skill service workers in all experiments. We increase the minimum wage by 5.5%, a value chosen to ensure constant expected earnings for low skill service workers between 1997 and 2016.

Table 3 summarizes our results. The first column reports results for the case without any minimum wage, and coincides with the results in Sections 5.1 to 5.4. The second column shows that a 5.5% minimum wage increase moderately attenuates the employment gains from outsourcing from 14.9 p.p. to 13.8 p.p. as it raises the cost of labor. The outsourcing expenditure share remains virtually unchanged, because outsourcing depends on wages at goods producers relative to wages at contractors. By contrast, the minimum wage pushes up all wages in the economy. For comparison, the third column reports the impact of the minimum wage increase without outsourcing demand and supply shocks.

³⁵Autor et al. (2020) find within-firm declines of the labor share, but reallocation towards firms with low labor shares. In our environment, outsourcing has opposite implications and can have neutral effects on the labor share.

Table 3: Outsourcing versus minimum wage and employment policies.

	Outsourcing	Out. with min. wage	Min. wage
Minimum wage (%)		5.50	5.50
Low-skill service non-emp. rate (p.p.)	-14.9	-13.8	0.37
Average wage (%)			
Low skill service	-6.81	-5.07	1.42
Other skills	-2.13	-1.42	0.51
AKM correlation (p.p.)	23.3	22.9	-0.31
Worker expected earnings (%)			
Low skill service	-1.49	0.00	0.99
Other skills	-0.67	-0.09	0.33
Outsourcing (expenditure share, p.p.)	9.52	8.95	0.02
Value added (%)	6.09	6.51	0.15
TFP (%)	-1.76	-1.01	0.49
Profits (%)	8.51	8.42	-0.32
Total welfare (%)	1.57	2.13	0.29

Perhaps surprisingly, output increases by a modest but additional 0.41% under the minimum wage raise. This additional increase in output is due to a 0.77% relative improvement in aggregate TFP: the minimum wage increase alleviates well-known composition externalities whereby low productivity firms hire too many workers in equilibrium (Acemoglu, 2001, Bilal, 2021). The reduction in these composition externalities interacts with outsourcing, as TFP gains surpass those with only the minimum wage increase and no outsourcing in the third column. Outsourcing exacerbates composition externalities, since contractors operate at the bottom of the job ladder and also post excess vacancies.

Profits increase by less with the minimum wage hike, but moderately enough that aggregate welfare rises. Overall, we conclude that a moderate increase in the minimum wage can ensure that outsourcing benefits workers through both employment and wage gains, at moderate costs for the shareholders of firms.

Conclusion

This paper started with a theory of domestic outsourcing. We have argued that it is useful to conceptualize the outsourcing decision of firms in the context of frictional labor markets. Monopsony power and firm wage premia emerge in equilibrium. More productive firms are more likely to outsource. Outsourcing raises output at the firm level. Contractors endogenously locate at the bottom of the job ladder, implying that outsourced workers receive lower wages. Together, these observations characterize the tension between productivity enhancements and redistribution away from workers

that is tied to outsourcing. Using firm-level instruments for outsourcing and revenue productivity, we have proposed new reduced-form evidence that confirms the productivity and redistributive effects of outsourcing. Finally, equipped with a structurally estimated model, we have shown that outsourcing deteriorates labor market prospects for low skill service workers. Accompanied by a minimum wage increase, outsourcing can benefit both workers and firms.

There are at least four natural directions along which to expand this research agenda. First, identifying explicit make-or-buy frictions and integrating them with our labor market theory of outsourcing may lead to novel policy implications. Second, the comparative advantage and cost-saving views can in principle be contrasted with the data for high skill workers for whom an outsourcing wage premium may arise. Third, our environment with outsourcing could be adapted to study the implications of the gig economy for inequality and output. Fourth, due to its tractability under parsimonious assumptions, our framework is naturally equipped to study questions with an efficiency-equity trade-off that involve wages and scale-biased aggregate transformations, such as trade liberalizations, automation or the rise of artificial intelligence.

References

- Abowd, John M., Kramarz, Francis, and Margolis, David N. (1999). “High Wage Workers and High Wage Firms”. *Econometrica* 67.2, pp. 251–333.
- Abraham, Katharine G. and Taylor, Susan K. (1996). “Firms’ Use of Outside Contractors: Theory and Evidence”. *Journal of Labor Economics* 14.3, pp. 394–424. (Visited on 10/24/2022).
- Acemoglu, Daron (2001). “Good Jobs versus Bad Jobs”. *Journal of Labor Economics* 19.1, pp. 1–21.
- Acemoglu, Daron, Gancia, Gino, and Zilibotti, Fabrizio (2015). “Offshoring and Directed Technical Change”. *American Economic Journal: Macroeconomics* 7.3, pp. 84–122.
- Antràs, Pol, Fort, Teresa C., and Tintelnot, Felix (2017). “The Margins of Global Sourcing: Theory and Evidence from U.S. Firms”. *American Economic Review* 107.9, pp. 2514–64.
- Antràs, Pol (Nov. 2003). “Firms, Contracts, and Trade Structure”. *The Quarterly Journal of Economics* 118.4, pp. 1375–1418.
- Autor, David, Dorn, David, Katz, Lawrence F, Patterson, Christina, and Van Reenen, John (Feb. 2020). “The Fall of the Labor Share and the Rise of Superstar Firms*[”]. *The Quarterly Journal of Economics* 135.2, pp. 645–709.
- Berg, Gerard J. van den and Ridder, Geert (1998). “An Empirical Equilibrium Search Model of the Labor Market”. *Econometrica* 66.5, pp. 1183–1221. (Visited on 04/18/2022).
- Bergeaud, Antonin, Malgouyres, Clement, Mazet-Sonilhac, Clement, and Signorelli, Sara (2020). “Technological Change and Domestic Outsourcing”. *Working Paper*.

- Berger, David, Herkenhoff, Kyle, and Mongey, Simon (2022). “Labor Market Power”. *American Economic Review* 112.4, pp. 1147–93.
- Berman, Abraham and Plemmons, Robert (1979). *Nonnegative Matrices in the Mathematical Sciences*. Academic Press.
- Bernhardt, Annette, Batt, Rosemary, Houseman, Susan N., and Appelbaum, Eileen (2016). “Domestic Outsourcing in the United States: A Research Agenda to Assess Trends and Effects on Job Quality”. *Upjohn Institute Working Paper*.
- Bertrand, Marianne, Hsieh, Chang-Tai, and Tsivanidis, Nick (2020). “Contract Labor and Firm Growth in India”. *NBER Working Paper 29151*.
- Bilal, Adrien (2021). “The Geography of Unemployment”. *NBER Working Paper 29269*.
- Bilal, Adrien G., Engbom, Niklas, Mongey, Simon, and Violante, Giovanni L. (2022). “Firm and Worker Dynamics in a Frictional Labor Market”. *Econometrica* 90.4, pp. 1425–1462.
- Bonhomme, Stéphane, Lamadon, Thibaut, and Manresa, Elena (2019). “A Distributional Framework for Matched Employer Employee Data”. *Econometrica* 87.3, pp. 699–739.
- Bontemps, Christian, Robin, Jean-Marc, and Berg, Gerard J. van den (2000). “Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation”. *International Economic Review* 41.2, pp. 305–358. (Visited on 04/18/2022).
- Borusyak, Kirill, Hull, Peter, and Jaravel, Xavier (June 2021). “Quasi-Experimental Shift-Share Research Designs”. *The Review of Economic Studies* 89.1, pp. 181–213.
- Bostanci, Gorkem (2021). “Productivity Gains from Labor Outsourcing: The Role of Trade Secrets”. *Working Paper*.
- Breza, Emily, Kaur, Supreet, and Shamdasani, Yogita (Oct. 2017). “The Morale Effects of Pay Inequality*”. *The Quarterly Journal of Economics* 133.2, pp. 611–663.
- Burdakov, Oleg, Dai, Yu-Hong, and Huang, Na (2019). “Stabilized Barzilai-Borwein method”. *Working Paper*.
- Burdett, Kenneth and Mortensen, Dale T. (1998). “Wage Differentials, Employer Size, and Unemployment”. *International Economic Review* 39.2, pp. 257–273.
- Card, David, Cardoso, Ana Rute, Heining, Joerg, and Kline, Patrick (2018). “Firms and Labor Market Inequality: Evidence and Some Theory”. *Journal of Labor Economics* 36.S1, S13–S70.
- Card, David, Heining, Jörg, and Kline, Patrick (2013). “Workplace Heterogeneity and the Rise of West German Wage Inequality”. *The Quarterly Journal of Economics* 128.3, pp. 967–1015.
- Card, David, Mas, Alexandre, Moretti, Enrico, and Saez, Emmanuel (2012). “Inequality at Work: The Effect of Peer Salaries on Job Satisfaction”. *American Economic Review* 102.6, pp. 2981–3003.

- Dorn, David, Schmieder, Johannes, and Spletzer, James (2018). “Domestic Outsourcing in the United States”. *Working Paper*.
- Drenik, Andres, Jäger, Simon, Plotkin, Miguel Pascuel, and Schoefer, Benjamin (2020). *Paying Outsourced Labor: Direct Evidence from Linked Temp Agency-Worker-Client Data*. Working Paper 26891. National Bureau of Economic Research.
- Dube, Arindrajit and Kaplan, Ethan (2010). “Does Outsourcing Reduce Wages in the Low-Wage Service Occupations? Evidence from Janitors and Guards”. *ILR Review* 63.2, pp. 287–306.
- Eaton, Jonathan and Kortum, Samuel (2002). “Technology, Geography, and Trade”. *Econometrica* 70.5, pp. 1741–1779. (Visited on 05/11/2022).
- Engbom, Niklas and Moser, Christian (2021). “Earnings Inequality and the Minimum Wage: Evidence from Brazil”. *NBER Working Paper 28831*.
- Feenstra, Robert and Hanson, Gordon (1999). “The Impact of Outsourcing and High-Technology Capital on Wages: Estimates For the United States, 1979–1990”. *The Quarterly Journal of Economics* 114.3, pp. 907–940.
- Felix, Mayara and Wong, Michael B. (2021). “Labor Market Consequences of Domestic Outsourcing: Evidence from Legalization in Brazil”. *Working Paper*.
- Giannoni, Marc and Mertens, Karel (2019). “Outsourcing, Markups and the Labor Share”. *Working Paper*.
- Goldschmidt, Deborah and Schmieder, Johannes F. (2017). “The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure”. *The Quarterly Journal of Economics* 132.3, pp. 1165–1217.
- Goldsmith-Pinkham, Paul, Sorkin, Isaac, and Swift, Henry (2020). “Bartik Instruments: What, When, Why, and How”. *American Economic Review* 110.8, pp. 2586–2624.
- Gouin-Bonenfant, Emilien (2022). “Productivity Dispersion, Between-firm Competition and the Labor Share”. *Forthcoming, Econometrica*.
- Grossman, Gene M. and Helpman, Elhanan (2002). “Integration versus Outsourcing in Industry Equilibrium”. *The Quarterly Journal of Economics* 117.1, pp. 85–120. (Visited on 10/14/2022).
- (2005). “Outsourcing in a Global Economy”. *The Review of Economic Studies* 72.1, pp. 135–159. (Visited on 04/18/2022).
- Grossman, Gene M. and Rossi-Hansberg, Esteban (2008). “Trading Tasks: A Simple Theory of Offshoring”. *American Economic Review* 98.5, pp. 1978–97.
- Grossman, Sanford J. and Hart, Oliver D. (1986). “Integration versus Outsourcing in Industry Equilibrium”. *Journal of Political Economy* 94.4, pp. 691–719.

- Haltiwanger, John C., Hyatt, Henry R., Kahn, Lisa B., and McEntarfer, Erika (2018). “Cyclical Job Ladders by Firm Size and Firm Wage”. *American Economic Journal: Macroeconomics* 10.2, pp. 52–85.
- Handwerker, Elizabeth W. (2021). “Outsourcing, Occupationally Homogeneous Employers, and Growing Wage Inequality in the United States”. *Working Paper*.
- Hart, Oliver and Moore, John (1990). “Property Rights and the Nature of the Firm”. *Journal of Political Economy* 98.6, pp. 1119–1158.
- Heise, Sebastian and Porzio, Tommaso (2022). “Labor Misallocation Across Firms and Regions”. *National Bureau of Economic Research Working Paper Series 30298*.
- Hummels, David, Jorgensen, Rasmus, Munch, Jakob, and Xiang, Chong (2014). “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data”. *American Economic Review* 104.6, pp. 1597–1629.
- Katz, Lawrence F. and Krueger, Alan B. (2017). “The Role of Unemployment in the Rise in Alternative Work Arrangements”. *American Economic Review* 107.5, pp. 388–92.
- (2019). “The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015”. *ILR Review* 72.2, pp. 382–416.
- Lamadon, Thibaut, Mogstad, Magne, and Setzler, Bradley (2022). “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market”. *American Economic Review* 112.1, pp. 169–212.
- LeMoigne, Mathilde (2020). “Exploring the ‘Fissured workplace’: Internal job ladders’ fragmentation and its effect on plants and workers”. *Working Paper*.
- Melitz, Marc J. (2003). “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity”. *Econometrica* 71.6, pp. 1695–1725.
- Mortensen, Dale T. and Vishwanath, Tara (1991). “Information Sources and Equilibrium Wage Outcomes”. *Northwestern University, Kellogg School of Management Discussion Paper 948*.
- Munoz, Mathilde (2022). “Trading Non-Tradables: The Implications of Europe’s Job Posting Policy”. *Working Paper*.
- Petrongolo, Barbara and Pissarides, Christopher A. (2001). “Looking into the Black Box: A Survey of the Matching Function”. *Journal of Economic Literature* 39.2, pp. 390–431.
- Segal, Lewis M. and Sullivan, Daniel G. (1997). “The Growth of Temporary Services Work”. *Journal of Economic Perspectives* 11.2, pp. 117–136.
- Song, Jae, Price, David J, Guvenen, Fatih, Bloom, Nicholas, and Wachter, Till von (Oct. 2018). “Firming Up Inequality*”. *The Quarterly Journal of Economics* 134.1, pp. 1–50.

- Sorkin, Isaac (Jan. 2018). “Ranking Firms Using Revealed Preference*”. *The Quarterly Journal of Economics* 133.3, pp. 1331–1393.
- Spitze, Scott (2022). “The Equilibrium Effects of Domestic Outsourcing”. *Working Paper*.
- Stock, James and Yogo, Motohiro (2005). “Testing for Weak Instruments in Linear IV Regression”. *Andrews DWK Identification and Inference for Econometric Models*, pp. 80–108.
- Topkis, Donald M. (1998). *Supermodularity and Complementarity*. Princeton University Press. ISBN: 9780691032443.

APPENDIX

A Proofs

A.1 Workers

We relate the wage offer distribution $F_s(w)$ to the wage distribution of employed workers $G_s(w)$ using worker flows. The flow of workers out of any wage interval $[\underline{w}_s, w)$ equals the flow of workers into that wage interval: $\lambda_s^U F_s(w) u_s = (\delta_s + \lambda_s^E (1 - F_s(w))) (m_s - u_s) G_s(w)$, where u_s denotes the skill-specific unemployment rate. The left-hand-side is the flow of workers out of unemployment into the wage interval $[\underline{w}_s, w)$, while the right-hand-side is the flow of workers out of that wage interval. It consists of workers who exogenously lose their job, and those who transition into higher wages. A similar argument guarantees that $u_s = \frac{m_s \delta_s}{\delta_s + \lambda_s^U}$. Re-arranging delivers (1).

$N_s(w)$ is equal to the limit of the ratio $\frac{G_s(w) - G_s(w-\varepsilon)}{F_s(w) - F_s(w-\varepsilon)}$ when $\varepsilon \rightarrow 0$, times the number of employed workers $m_s - u_s$. Straightforward differentiation delivers (2).

A.2 Reservation wage

Omit s indices. Suppose without loss of generality that F admits a density f . Then $\left[r + \delta + \lambda^E (1 - F(w)) \right] V(w) = w + \delta U + \lambda^E \int_w^\infty V(x) f(x) dx$. Differentiate w.r.t. w to obtain $\left[r + \delta + \lambda^E (1 - F(w)) \right] V'(w) = 1$. Integrate back to $V(w) = U + \int_w^\infty \frac{dx}{r + \delta + \lambda^E (1 - F(x))}$. Substituting into the value of unemployment, $rU = b + \lambda^U \int_w^\infty \frac{(1 - F(x)) dx}{r + \delta + \lambda^E (1 - F(x))}$. Since $V(\underline{w}) = U$, $(r + \lambda^U)U = b + \lambda^U \int_w^\infty V(x) f(x) dx$ and $(r + \lambda^E)U = \underline{w} + \lambda^E \int_w^\infty V(x) f(x) dx$. Thus, $rU = \frac{\lambda^U \underline{w} - \lambda^E b}{\lambda^U - \lambda^E}$. Therefore,

$$\lambda^U \underline{w} = \lambda^E b + (\lambda^U - \lambda^E) \left[b + \lambda^U \int_w^\infty \frac{(1 - F(x)) dx}{r + \delta + \lambda^E (1 - F(x))} \right] \quad (26)$$

A.3 Proof of Proposition 1

Impose Assumption (B). Then $n_s(w) = \frac{n_{0s}}{[1 + k_s(1 - F_s(w))]^2}$, and $n'_s(w) = \frac{2kn_{0s}F'_s(w)}{[1 + k_s(1 - F_s(w))]^3}$. (3) becomes

$$\pi(z) = \max_{v_s \in [0, 1]^S, w_s} R(z, \{n_s(w)v_s\}_s) - \int w_s n_s(w_s) v_s ds. \quad (27)$$

Start from the FOC for wages in (3). We obtain $R_{n_s} n'_s - n_s - w_s n'_s = 0$.³⁶ Differentiating the objective in (27) w.r.t. v_s and using the FOC for wages, we obtain $\frac{\partial(R(z, \{n_s(w)v_s\}_s) - \int w_s n_s(w_s) v_s ds)}{\partial v_s} = (R_{n_s} - w_s) n_s > 0$. Thus, firms are always at the corner $v_s = 1$. Hence, (3) coincides with

$$\pi(z) = \max_{w_s} \Pi[z, \{w_s\}_s] \equiv R(z, \{n_s(w)\}_s) - \int w_s n_s(w_s) ds. \quad (28)$$

³⁶This equality implies $(R_{n_s} - w_s) n'_s = n_s > 0$. Thus, $R_{n_s} > w_s$ and $n'_s > 0$.

Since n_s is increasing in w , Π is continuously differentiable and strictly supermodular in any pair (z, w_s) . In addition, the profit function is supermodular in $\{w_s\}_s$, and exhibits increasing differences in (z, w_s) for all s . In addition, the set of $\{w_s\}_s$ forms a lattice with the element-wise order. Therefore, we can apply Theorem 2.8.5. p. 79 in Topkis (1998). Thus, the set of maximizers $\{w_s(z)\}_s$ are strictly increasing in z for each s . Given the ordering of the ordering of wages, $F(w_s(z)) = \Gamma(z)$

A.4 Proof of Proposition 2

Given Proposition 1, it is immediate to verify that $n_s(z) = \frac{(1+k_s)e_s}{M[1+k_s(1-\Gamma(z))]}$.

A.5 Proof of Proposition 3

Because wages are strictly increasing in z , they are continuous almost everywhere and we may take first-order conditions for almost every productivity z . Hence:

$$\frac{d(n_s(w)w)}{dw} \Big|_{w=w_s(z)} = \frac{dR(z, \mathbf{n}_{-s}(w_s(z)), n_s(w))}{dw} \Big|_{w=w_s(z)} = \frac{\partial R}{\partial n_s}(z, \mathbf{n}(z)) \cdot n'_s(w_s(z))$$

where \mathbf{n}_{-s} denotes the vector \mathbf{n} without its entry s . Multiplying both sides by $w'_s(z)$ and changing variables to $n_s(w_s(z)) \equiv n_s(z)$ delivers

$$n_s(z)w'_s(z) = n'_s(z)(R_{n_s}(z, \{n_t(z)\}_t) - w_s(z)). \quad (29)$$

Integrating over z subject to the boundary condition $w_s(z) = \underline{w}_s$ delivers the formula in Proposition 3.

A.6 Proof of Proposition 4

Existence and uniqueness among equilibria with continuous F . Together, Propositions 1, 2 and 3 suffice to complete a guess and verify strategy to exhibit an equilibrium with a continuous wage offer distribution. The last condition to verify is whether a reservation wage compatible with those results exists. Omitting s subscripts, re-write (26) as

$$\lambda^U \underline{w} = \lambda^E b + (\lambda^U - \lambda^E) \left[b + \lambda^U \int_{\underline{w}}^{\infty} \frac{(1 - \Gamma(x))w'(x)dx}{r + \delta + \lambda^E(1 - \Gamma(x))} \right]. \quad (30)$$

$w'(x)$ is a function of the reservation wage \underline{w} through the ODE (29). To explicit its dependence, denote $d(z) = \frac{\partial w(z)}{\partial \underline{w}}$ where the partial derivative is understood as a derivative w.r.t. the initial condition of the ODE (29). Differentiating (29), we obtain $n(z)d'(z) = -n'(z)d(z)$. Solving this ODE explicitly and using $d(\underline{z}) = 1$ by definition, we obtain $d(z) = \frac{n(\underline{z})}{n(z)}$. Hence, $d'(z) = -\frac{n(\underline{z})n'(z)}{n(z)^2} < 0$. Thus, $w'(x)$ is a decreasing function of \underline{w} .

Hence, the right-hand-side of (30) is a decreasing function of \underline{w} that goes to $\lambda^E b \leq \lambda^U b$ as \underline{w} goes to infinity. Its left-hand-side is an increasing function of \underline{w} that spans $\lambda^U b$ to $+\infty$. Therefore, there exists a unique reservation wage \underline{w} .

Existence and uniqueness among all possible equilibria. For expositional simplicity and without loss of generality, we present our trembling-hand refinement with a single skill. The maximization problem (3) becomes:

$$w(z) = \operatorname{argmax}_{w,n} R(z, n) - wn, \quad n \leq n(w) \equiv \frac{n_0}{[1 + k(1 - F(w))][1 + k(1 - F(w^-))]}.$$

Suppose that firms make mistakes ε after choosing their target wage: firms post $w(z) + \varepsilon$ while having chosen $w(z)$ for an i.i.d. shock ε across firms. Firm z does not expect to make a mistake, but takes into account the equilibrium wage offer distribution inclusive of other firms' mistakes.

The distribution F that enters the constraint is the vacancy-weighted distribution of posted wages $w + \varepsilon$ in the economy. We impose the following assumptions on the distribution of mistakes ε , H_σ . H_σ has a \mathcal{C}^∞ density with compact support. The variance of H_σ (or any relevant measure of dispersion such as the size of its support) is given by σ . Convergence of H_σ is uniform: $|H_\sigma(\varepsilon) - \mathbf{1}\{\varepsilon \geq 0\}| \leq h_0\sigma$, where h_0 is a constant independent from σ . H_σ is strictly increasing. Standard results on convolution kernels imply that there exists such a distribution. Our trembling-hand refinement is to consider the limiting economy when $\sigma \downarrow 0$.

Assumption (B). When $\sigma = 0$, we restrict attention to decentralized equilibria that are the limit of a sequence of decentralized equilibria when $\sigma \downarrow 0$.

The remainder of this section is structured as follows. First, we show that the wage distribution is smooth when $\sigma > 0$. Second, we show that wages are strictly increasing in z when $\sigma > 0$. Third, we show that the wage rank converges to the productivity rank as $\sigma \rightarrow 0$. Fourth, we show that wages converge to our candidate equilibrium when $\sigma \rightarrow 0$.

1. Smooth wage distribution when $\sigma > 0$. F is a convolution between the distribution of chosen wages $w(z)$ and an i.i.d. shock ε . Therefore, standard results on regularizing convolutions ensure that F admits a \mathcal{C}^∞ density when $\sigma > 0$. This conclusion follows from $F(w) = \int H_\sigma(w - \omega) d\Omega(\omega)$ together with dominated convergence, where $\omega = w(z)$ is a random variable that denotes chosen wages, and Ω is its c.d.f. In addition, F is strictly increasing: $F'(w) > 0$. Since F is smooth, for any $\sigma > 0$, $n(w) = \frac{n_0}{[1+k(1-F(w))]^2}$ and $n'(w) = \frac{2kn_0F'(w)}{[1+k(1-F(w))]^3}$.

2. Binding constraint and increasing wages when $\sigma > 0$. Conditional on F being smooth, the argument is identical to Section A.3. Crucially, this conclusion would not be valid in general if there were a mass point in the distribution F .

3. Wage rank and productivity rank when $\sigma \downarrow 0$. Denote by $w_\sigma(z)$ the wage function for a given σ . Write $F(w(z_0)) = \mathbb{P}[\varepsilon \leq w(z_0) - w(z)] = \int H_\sigma(w(z_0) - w(z)) d\Gamma(z)$. Then

$$\int H_\sigma(w_\sigma(z_0) - w_\sigma(z)) d\Gamma(z) = \underbrace{\int \mathbf{1}\{z \leq z_0\} d\Gamma(z)}_{\text{By wage ranking}} + \underbrace{\int [H_\sigma(w_\sigma(z_0) - w_\sigma(z)) - \mathbf{1}\{w_\sigma(z_0) - w_\sigma(z) \geq 0\}]}_{\leq h_0\sigma \text{ by assumption}} d\Gamma(z)$$

Therefore, for all z , $F(w(z)) \rightarrow \Gamma(z)$ uniformly, and $n(w(z)) \rightarrow n(z) \equiv \frac{n_0}{[1+k(1-\Gamma(z))]^2}$ uniformly.

4. Wages when $\sigma \downarrow 0$. We go back to the maximization problem (27) and use an argument that resembles Berge's maximum theorem that we cannot apply directly. Re-write (27) as choosing the wage $w_\sigma(Z)$ of a firm with productivity Z . The wage function w_σ must hence satisfy

$z = \operatorname{argmax}_Z R(z, n(w_\sigma(Z))) - w_\sigma(Z)n(w_\sigma(Z))]$. In particular, $Z^*(z) = z$ for all σ . Suppose for a contradiction that w_σ was discontinuous in σ at $\sigma = 0$ for some z_0 . Since $n(w_\sigma(Z)) \rightarrow n(Z)$, it must be that $Z^*(z)$ jumps down at $\sigma = 0$ since firms downscale due to higher costs of labor. This contradicts $Z^*(z_0) = z_0$. Therefore, w_σ is continuous in σ at $\sigma = 0$. At $\sigma = 0$, w_0 satisfies $z = \operatorname{argmax}_Z R(z, n(Z)) - w_0(Z)n(Z)$. w_0 thus solves $(R_n(z, n(z)) - w_0(z))n'(z) = w_0'(z)n(z)$, which coincides with the wage ODE in our candidate equilibrium. Thus, the limit of any equilibrium under Assumption (B) as $\sigma \downarrow 0$ converges to the candidate equilibrium.

A.7 Micro-foundations for the cost of outsourcing

Iceberg trade cost or productivity wedge. To sell one unit of labor services to a goods producer, contractor firms must hire $1/\tau_s$ units of labor.

Capital. Assume that contractor firms for skill s combine capital, in exogenous supply K_s , and labor to produce one unit of efficiency unit of labor services of a given skill s . The decision problem of the contractor firm is

$$\pi^C(w) = \max_k p_s k^{1-\beta} n_s(w)^\beta - r_s k - w n_s(w). \quad (31)$$

The optimality condition for capital is then $k = \left(\frac{(1-\beta)p_s}{r_s}\right)^{\frac{1}{\beta}} \cdot n_s(w)$. Market clearing for capital leads to $\frac{r_s}{1-\beta} = p_s(Q_s^{\text{Out}}/K_s)^\beta$ where Q_s^{Out} is aggregate employment in contractor firms. Substituting back into (31), we obtain $\pi^C(w) = p_s \left(\frac{K_s}{Q_s^{\text{Out}}}\right)^{1-\beta} n_s(w) - w n_s(w)$. Assume further that $K_s = \tau_s^{\frac{1}{1-\beta}}$, and take $\beta \rightarrow 1$. Then, (31) becomes $\pi^C(w) = (\tau_s p_s - w)n_s(w)$.

A.8 Proof of Proposition 5

We start by defining the cost function

$$C_s(n) = \min \{w_s(n)n, p_s n\} = n \min \{w_s(n), p_s\},$$

where $w_s(n)$ is the inverse function of $n_s(w)$. Then rewrite the profit-maximization problem (5) as

$$\pi(z) = \max_{\{n_s\}_s} R(z, \{n_s\}_s) - \int C_s(n_s) ds. \quad (32)$$

As in Proposition 1, the profit function in (32) is supermodular in $(z, \{n_s\}_s)$. We again use Theorem 2.8.1. p. 76 in Topkis (1998) to obtain that size is rising in productivity: $n_s^*(z)$ is increasing in z for every s . Given that w_s is increasing in n and $n_s^*(z)$ is increasing in z , there must exist a threshold \hat{z}_s such that the minimum of the cost function is attained in-house for $z \leq \hat{z}_s$, and attained outsourced for $z > \hat{z}_s$.

A.9 Proof of Proposition 6

For clarity of exposition, we first present our proof when there is a single worker type, before moving to the multiple worker types case.

A.9.1 Single worker type

Focus on the case with a single worker type in the remainder of this proof and drop the s index. We denote effective employment by n , and sometimes write n_{in} or n_{out} to denote in-house vs. outsourced employment.

Proof that $n_{\text{in}}(\hat{z}) < n_{\text{out}}(\hat{z})$. The theorem of the maximum implies that profits are continuous at the outsourcing cutoff \hat{z} . The indifference condition at \hat{z} writes

$$R(\hat{z}, n_{\text{in}}(\hat{z})) - w(n_{\text{in}}(\hat{z}))n_{\text{in}}(\hat{z}) = R(\hat{z}, n_{\text{out}}(\hat{z})) - pn_{\text{out}}(\hat{z}) \quad (33)$$

where $w(n)$ is the (increasing) inverse function of the labor supply curve. The first-order condition for in-house employment is $R_n(\hat{z}, n_{\text{in}}(\hat{z})) = w(\hat{z}) + n_{\text{in}}(\hat{z})w'(n_{\text{in}}(\hat{z}))$. For outsourcing it is $R_n(\hat{z}, n_{\text{out}}(\hat{z})) = p$. Substituting both into the indifference condition (33):

$$R(\hat{z}, n_{\text{in}}(\hat{z})) - R_n(\hat{z}, n_{\text{in}}(\hat{z}))n_{\text{in}}(\hat{z}) + n_{\text{in}}(\hat{z})^2w'(n_{\text{in}}(\hat{z})) = R(\hat{z}, n_{\text{out}}(\hat{z})) - n_{\text{out}}(\hat{z})R_n(\hat{z}, n_{\text{out}}(\hat{z}))$$

R being strictly concave in n , the function $R - nR_n$ is strictly increasing in n . Since $w'(n(\hat{z})) > 0$, we immediately obtain that $n_{\text{in}}(\hat{z}) > n_{\text{out}}(\hat{z})$.

Proof that $R(\hat{z}, n_{\text{in}}(\hat{z})) < R(\hat{z}, n_{\text{out}}(\hat{z}))$. Since effective employment jumps up at the outsourcing cutoff, so does revenue.

A.9.2 Multiple worker types

We now return to the case with multiple worker types. We consider what happens to employment n_s of a particular skill s when a firm outsources. The overall structure of the proof follows closely that with a single worker type. The only difference is that we must control the difference in labor costs for other worker types $t \neq s$ in the indifference condition.

For notational simplicity, we omit dependence on \hat{z} in this subsection since all function are evaluated at this value. We also denote for instance $R_{\text{in}} \equiv R(\hat{z}, n_{s,\text{in}}(\hat{z}), n_{-s,\text{in}})$ revenues at the outsourcing threshold when skill s is in-house, and employment at all skills $t \neq s$ are chosen optimally given the in-house choice for s .

Proof that $n_{s,\text{in}}(\hat{z}) < n_{s,\text{out}}(\hat{z})$. The indifference condition between in-house and outsourcing becomes

$$R_{\text{in}} - n_{s,\text{in}}R_{n_{s,\text{in}}} + (n_{s,\text{in}}^{\text{in}})^2w'_s(n_{s,\text{in}}) - C_{-s}^{\text{in}} = R_{\text{out}} - n_{s,\text{out}}R_{n_{s,\text{out}}} - C_{-s}^{\text{out}} \quad (34)$$

where C_{-s} denotes the cost function for skills other than s . Our goal is to show that $n_s^{\text{out}} > n_s^{\text{in}}$. We now guess and verify that $n_s^{\text{out}} > n_s^{\text{in}}$.

We seek to show that, under the guess, $C_{-s,\text{out}} - C_{-s,\text{in}} > 0$. Fix a particular $t \neq s$. Write the F.O.C.s for n_t depending on whether skill s is in-house or outsourced as

$$R_{n_t}(n_{s,\text{in}}, n_{-s,\text{in}}) = C_{t,n_t}(n_{t,\text{in}}) \quad , \quad R_{n_t}(n_{s,\text{out}}, n_{-s,\text{out}}) = C_{t,n_t}(n_{t,\text{out}})$$

Thus, we only need to study how $n_t(x)$ varies with x in the equation indexed by x :

$$R_{n_t}(x, n_{-s}(x)) = C_{n_t}(n_t(x)). \quad (35)$$

Differentiating (35) w.r.t. x , we obtain for all $t \neq s$:

$$R_{n_s n_t} + \sum_{k \neq s} R_{n_t n_k} n'_k(x) = C_{n_t n_t} n'_t(x) \implies R_{n_s n_t} + \sum_{k \neq s, t} R_{n_t n_k} n'_k(x) = (C_{n_t n_t} - R_{n_t n_t}) n'_t(x). \quad (36)$$

Define the matrix

$$A_{tk} = \begin{cases} C_{n_t n_t} - R_{n_t n_t} > 0 & \text{if } t = k \\ -R_{n_t n_k} < 0 & \text{if } t \neq k \end{cases}$$

where the inequality in the first row follows from the S.O.C. and the inequality in the second row from supermodularity. We thus re-write our problem in vector form

$$A \cdot \{n'_\ell\}_{\ell \neq s} = \{R_{n_s n_\ell}\}_{\ell \neq s} > 0$$

Importantly, A_{tk} is an M-matrix. As a result A^{-1} has positive elements (see e.g. Berman and Plemmons, 1979). We conclude that $n'_\ell > 0$ for $\ell \neq s$, and so

$$n_{-s,\text{out}} > n_{-s,\text{in}} \quad (37)$$

Equation (37) ensures that size of other skills increase when firms outsource skill s . Because costs rise with size, we obtain that

$$C_{-s}^{\text{out}} - C_{-s}^{\text{in}} > 0$$

We can now go back to the indifference condition (34). We thus have guaranteed that

$$(n_s^{\text{in}})^2 w'_s(n_s^{\text{in}}) + C_{-s}^{\text{out}} - C_{-s}^{\text{in}} > 0$$

As in the one-skill case, $R - n_s R_{n_s}$ is increasing because of concavity, and so we verify our guess that $n_{s,\text{out}} > n_{s,\text{in}}$.

Proof that $R(\hat{z}, n_{s,\text{in}}(\hat{z}), n_{-s,\text{in}}(\hat{z})) < R(\hat{z}, n_{s,\text{out}}(\hat{z}), n_{-s,\text{out}}(\hat{z}))$. This inequality immediately follows from $n_{s,\text{out}} > n_{s,\text{in}}$ together with $n_{-s,\text{out}} > n_{-s,\text{in}}$ and R being increasing in each argument.

A.10 Proof of Proposition 7

We start our proof with the case in which contractors have neutral or low advantage $\tau_s \leq 1$. We then turn to the case of comparative advantage $\tau_s \gg 1$.

A.10.1 Neutral or comparative disadvantage case $\tau_s \leq 1$

Proof that $w_s^{\text{cont.}} < w_s(\hat{z}_s)$. The F.O.C. for the marginal outsourcer at \hat{z}_s writes

$$p_s = R_{n_s}(\hat{z}_s, n_{s,\text{out}}(\hat{z}_s), \mathbf{n}_{-s,\text{out}}(\hat{z}_s)) \quad (38)$$

Suppose for a contradiction that $w_{s,\text{in}}(\hat{z})$ is in the support of the wage distribution of contractors. Then the wage F.O.C. holds for goods producer \hat{z} at $w = w_{s,\text{in}}(\hat{z})$. Then, the marginal benefit from a wage change dw for a contractor at $w = w_{s,\text{in}}(\hat{z})$ is, up to $n'_s(w)$, propotional to

$$\begin{aligned} & \left(p_s \tau_s - w - \frac{n_s(w)}{n'_s(w)} \right) \Big|_{w=w_{s,\text{in}}(\hat{z})} dw \\ \leq & \left(p_s - w - \frac{n_s(w)}{n'_s(w)} \right) \Big|_{w=w_{s,\text{in}}(\hat{z})} dw && \text{(use } \tau_s \leq 1) \\ = & \left(R_{n_s}(\hat{z}_s, n_{s,\text{out}}(\hat{z}), \mathbf{n}_{-s,\text{out}}(\hat{z})) - w - \frac{n_s(w)}{n'_s(w)} \right) \Big|_{w=w_{s,\text{in}}(\hat{z})} dw && \text{(use (38))} \\ < & \left(R_{n_s}(\hat{z}_s, n_{s,\text{in}}(\hat{z}), \mathbf{n}_{-s,\text{out}}(\hat{z})) - w - \frac{n_s(w)}{n'_s(w)} \right) \Big|_{w=w_{s,\text{in}}(\hat{z})} dw && \text{(use } n_{s,\text{in}}(\hat{z}) < n_{s,\text{out}}(\hat{z})) \end{aligned} \quad (39)$$

Our last step is to bound the marginal product $R_{n_s}(\hat{z}_s, n_{s,\text{in}}(\hat{z}), \mathbf{n}_{-s,\text{out}}(\hat{z}))$ of a firm that chooses to hire skill s in-house but holding its employment fixed at the optimal size when skill s is outsourced. To do so, we define a modified profit function

$$P(n; m_{-s}) = R(\hat{z}_s, n, m_{-s}) - w_s(n)n - C_{-s}(m_{-s})$$

where we hold productivity fixed at \hat{z} and m_{-s} fixed too. When evaluated at $m_{-s} = n_{s,\text{out}}$, this modified profit function corresponds to the firm considering achieving the size $n_{s,\text{out}}$ and modified outsourced profits, but by hiring skill s in-house. When evaluated at $m_{-s} = n_{s,\text{in}}$, this profit function is simply the fully in-house profit function.

Supermodularity implies that $\text{argmax}_n P(n, m_{-s})$ is increasing in m_{-s} . Since Proposition 6 implies $n_{-s,\text{out}} > n_{-s,\text{in}}$, we obtain

$$\text{argmax}_n P(n, n_{-s,\text{out}}) > \text{argmax}_n P(n, n_{-s,\text{in}}) = n_{s,\text{in}}$$

As a result, $n_{s,\text{in}}$ lies to the right of the maximizer of $n \mapsto P(n, n_{-s,\text{out}})$, and so $n \mapsto P(n, n_{-s,\text{out}})$ must be decreasing when evaluated at $n = n_{s,\text{in}}$:

$$P_n(n_{s,\text{in}}, n_{-s,\text{out}}) < 0$$

This last condition rewrites

$$R_n(n_{s,\text{in}}, n_{-s,\text{out}}) < \frac{\partial}{\partial n} (w_s(n)n) \Big|_{n=n_{s,\text{in}}} \quad (40)$$

Substituting (40) into the last line of (39), we finally obtain

$$\left(p_s \tau_s - w - \frac{n_s(w)}{n'_s(w)} \right) \Big|_{w_s=w_{s,\text{in}}(\hat{z})} dw < 0 \quad (41)$$

Thus, contractors who consider posting wage $w_{s,\text{in}}(\hat{z})$ prefer to lower their wage offer. This observation implies that contractors post wages $w_s^{\text{cont.}}$ strictly below $w_{s,\text{in}}(\hat{z})$.

Proof that $w_s^{\text{cont.}} \equiv \underline{w}_s$ under free-entry. Impose free-entry. If contractors posted any wage $w > \underline{w}_s$, they could deviate to $w - \varepsilon$ for a small ε and make positive profits, a contradiction. Thus, they all post the reservation wage.

A.10.2 Comparative advantage case $\tau_s \gg 1$

We start our proof under the assumption of a single worker type and an isoelastic revenue function. Then we relax the assumption of isoelastic revenues.

Isoelastic revenue function. Suppose that $R(z, n) = zn^\rho$ with $0 < \rho < 1$. Market clearing for labor services writes

$$M \int_{\hat{z}}^{\bar{z}} \left(\frac{\rho z}{p} \right)^{\frac{1}{1-\rho}} d\Gamma(z) = \tau \times S \int n_{\text{out}}(w) dF(w)$$

Since in-house employment is always bounded below and above, $S \int n_{\text{out}}(w) dF(w) \equiv N_{\text{out}}$ remains bounded. Hence, to a leading order when $\tau \rightarrow \infty$, market clearing implies

$$\int_{\hat{z}}^{\bar{z}} (\rho z)^{\frac{1}{1-\rho}} d\Gamma(z) \sim \tau p^{\frac{1}{1-\rho}}$$

\hat{z} is defined by the indifference condition $R - nR_n + n^2 w'(n) \Big|_{\text{in}} = R - nR_n \Big|_{\text{out}}$. In the isoelastic case it implies $c_1 \hat{z}^{\frac{1}{1-\rho}} p^{-\frac{\rho}{1-\rho}} \gtrsim c_2 \hat{z} + c_3$ for bounded and non-vanishing functions of τ , $c_i(\tau)$. So if $p \rightarrow 0$, this identity implies $\hat{z} \rightarrow \underline{z}$, all firms outsource, and the indifference condition ceases to hold.

Now we are ready for a guess and verify. We guess that, as $\tau \rightarrow +\infty$, $p \rightarrow 0$. Then \hat{z} hits \underline{z} . The integral in market clearing then becomes constant, and market clearing implies

$$1 \sim \tau p^{\frac{1}{1-\rho}} \implies p \sim \tau^{\rho-1} \rightarrow 0$$

and the guess is verified. This then proves that $\tau p \sim \tau^\rho \rightarrow +\infty$: the marginal product of labor of contractors becomes infinite. A similar wage deviation argument to the one in Section A.10.1 then ensures that goods producers post wages below contractors.

Relaxing isoelastic revenue function. As long as $0 < \underline{\rho} \leq \frac{nR_n}{R} \leq \bar{\rho} < 1$, the previous arguments continue to apply since they are all estimates that only rely on $0 < \rho < 1$.

A.11 Proof of Proposition 8

The hire rate from employment is increasing in w , while churn is decreasing in w . The comparative statics in Proposition 8 for these outcomes immediately follow from Proposition 7.

To show that net poaching is increasing in w as well, combine equations (1) and (2) to obtain $\frac{G_s(w)}{n_s(w)} = \frac{\widehat{M}_s}{(1+k_s)e_s} F_s(w)(1+k_s(1-F_s(w)))$. From Online Appendix F.1, $q_s = \frac{e_s \delta_s (1+k_s)}{\widehat{M}_s}$. Hence $\text{NP}_s(w) = (1-\phi_s)F_s(w)(\delta_s + \lambda_s^E(1-F_s(w))) - \lambda_s^E(1-F_s(w)) = (1-\phi_s)\delta_s F_s(w) - \lambda_s^E(1-F_s(w))[1 - (1-\phi_s)F_s(w)]$. Viewed as a function of $F_s(w)$, this quantity is weakly increasing in $F_s(w)$ if the second component $F \mapsto (1-F)[1 - (1-\phi)F]$ is decreasing. This component is a second-order polynomial with roots equal to 1 and $1/(1-\phi) > 1$, with a positive coefficient on the quadratic term. Hence, it is decreasing on $[0, 1]$. Hence, net poaching is increasing in w . We conclude the proof once more with Proposition 7.

A.12 Outsourcing equilibrium

With some outsourcing in equilibrium and in the weakly neutral case $\tau_s \leq 1$, the wage distribution for a given skill s has three regions. In the first, low wage region, only unproductive goods producers operate. In an intermediate wage region, contractors operate together with mid-productivity goods producers. In a high-wage region, only highly productive goods producers operate. Depending on parameter values, either the low-wage or the high-wage region may be empty. For the sake of brevity, we omit a lengthy distinction of all cases and describe the economy with all regions populated in the weakly neutral case $\tau_s \leq 1$. We denote by $\widehat{M}_s = M^C + M^G \Gamma(\hat{z}_s)$ the total measure of firms that hire skill s in-house.

Denote by z_{1s} the threshold productivity at which the low wage region ends, and z_{2s} the threshold productivity at which the high wage region starts. Goods producers $z \in [z, z_{1s}]$ behave similarly to the no-outsourcing economy. These goods producers are now poached by both contractors and high-productivity in-house firms. Their equilibrium rank in the job ladder and equilibrium size are

$$\Upsilon_{1s}(z) = \frac{M^G \Gamma(z)}{\widehat{M}_s}, \quad n_{1s}(z) = \frac{n_{0s}}{\widehat{M}_s [1 + k_s (1 - \Upsilon_{1s}(z))]^2}, \quad z \in [z, z_{1s}]. \quad (42)$$

where $n_{0s} = (1+k_s)e_s$. For goods producers in this low wage region, the only change to their wages relative to Proposition 3 stems from the number of workers they attract and retain with a wage offer given in equations (42). At the upper end of this low wage region, the marginal product of labor of goods producers equals that of the first contractor firm which determines the threshold z_{1s} : $R_{n_s}(z_{1s}, \mathbf{n}(z_{1s})) = p_s$.

In the intermediate wage region, contractors compete with in-house goods producers. Contractors being homogeneous, there are indifferent between paying any wage in this intermediate region. There, rank in the job ladder and size at any wage are directly determined by contractors' indifference

condition:

$$F_{2s}(w) = 1 + \frac{1}{k_s} \left[1 - \left(\frac{n_{0s}}{\widehat{M}_s n_{2s}(w)} \right)^{\frac{1}{2}} \right], \quad n_{2s}(w) = \frac{p_s - w(z_{1s})}{p - w} n_{1s}(z_{1s}), \quad z \in [z_{1s}, z_{2s}]. \quad (43)$$

Throughout this intermediate region, wages of goods producers keep rising with productivity so that the marginal product of labor of goods producers equals that of contractors: $R_{n_s}(z, n_{2s}(w_s(z)), n_{-s}^*(z)) = p_s$. The threshold productivity z_{2s} is reached when there are no more contractors left $F_{2s}(w_s(z_{2s})) = M^C + M^G \Gamma(z_{2s})$.

The economy then enters the third, high wage region with only highly productive goods producers. This region resembles the low wage region in that size and rank are given by

$$\Upsilon_{3s}(z) = \frac{M^C + M^G \Gamma(z)}{\widehat{M}_s}, \quad n_{3s}(z) = \frac{n_{0s}}{\widehat{M}_s [1 + k_s (1 - \Upsilon_{3s}(z))]^2}, \quad z \in [z_{2s}, \hat{z}_s]. \quad (44)$$

Starting from $w_s(z_{2s})$, wages once again follow Proposition 3 but with an equilibrium size given in equations (44). Above the threshold productivity \hat{z}_s , no firm operates in-house. There, firms outsource their employment.

A.13 Proof of Corollary 1

Write the problem of the firm as

$$\pi(z, \{\varepsilon_s\}_s) = \max_{n_s} R(z, \{n_s\}) - \sum_s C_s(n_s, \varepsilon_s) \quad , \quad C_s(n, \varepsilon) = n \min\{w_s(n), p_s \varepsilon\}$$

The profit function has the same supermodularity properties as before, but is also supermodular in $(n_s, 1/\varepsilon_s)$. Hence, optimal size $n_s^*(z, \varepsilon_s, \{\varepsilon_{s'}\}_{s' \neq s})$ is increasing in z and weakly decreasing in ε_s .

Evaluating the cost function C_s at optimal size n_s^* , one immediately obtains that there is a threshold productivity $\hat{z}_s(\{\varepsilon_{s'}\}_{s'})$ at which the firm switches between both parts of the cost function, when $p_s \leq w_s(n_s^*(\hat{z}_s(\{\varepsilon_{s'}\}_{s'}), \{\varepsilon_{s'}\}_{s'}))/\varepsilon_s$. The right-hand-side is decreasing in ε_s , and so the threshold $\hat{z}_s(\{\varepsilon_{s'}\}_{s'})$ is increasing in ε_s .

A.14 Proof of Corollary 2

The proof follows exactly that of Proposition 6 once we condition on the idiosyncratic shock ε_s .

A.15 Proof of Corollary 3

The proof follows exactly that of Proposition 7 once we condition on the shock ε_s and use $\varepsilon_s \leq 1$ in the series of inequalities (39). The proof for the comparative advantage case also remains unchanged because market clearing for outsourcing becomes

$$M^G \int d\Gamma(\boldsymbol{\varepsilon}) \int_{\hat{z}_s(\boldsymbol{\varepsilon})}^{\bar{z}} n_s^*(z, \boldsymbol{\varepsilon}) d\Gamma(z|\boldsymbol{\varepsilon}) = \tau_s M^C \int n_s^{\text{cont.}}(w) dF_s(w)$$

and thus leaves the proof in Appendix A.10 unchanged.

A.16 Proof that the outsourcing share rises with $1/\varepsilon_s$

Consider the setup from Section 4, and in particular the revenue function (18). The outsourcing share out of revenues conditional on outsourcing is $S(z, \boldsymbol{\varepsilon}) = \rho a_1$. The outsourcing share conditional on not outsourcing is 0. $\hat{z}_s(\boldsymbol{\varepsilon})$ is increasing in ε_1 (Corollary 1), so that the outsourcing share is weakly increasing in $1/\varepsilon_1$, and strictly increasing when $\hat{z}_s(\boldsymbol{\varepsilon})$ crosses z . Hence, $\mathbb{E} \left[\frac{\partial S(z, \boldsymbol{\varepsilon})}{\partial (1/\varepsilon_1)} \right] > 0$.

A.17 Welfare and expected earnings

Welfare. The value function of a worker with state $x \in \{b, w\}$, where b denotes unemployment, satisfies $V(x) = r \mathbb{E}_0 \int_0^\infty e^{-rt} x_t dt$. We rescale values by the discount rate r as we require values to remain finite in the limit $r \rightarrow 0$. We show that, when $r \rightarrow 0$, the value of any worker, regardless of their state, converges to steady-state expected earnings $\mathbb{E}[x]$ when the process x_t has a unique invariant distribution. Denote by $h(t, x)$ the solution to the time-dependent Kolmogorov Forward equation satisfied by the density of x_t . Then $V(x) = \int x r \int_0^\infty e^{-rt} h(t, x) d\mu(x)$ where μ is a base measure that has a Dirac mass point at b and is the Lebesgue measure for all $x \geq \underline{w}$. If we can show that $r \int_0^\infty e^{-rt} u(t) dt \rightarrow \lim_{t \rightarrow \infty} u(t)$ for any smooth and bounded function u , we can apply this last result x by x and obtain $V(x) \rightarrow \int x h(\infty, x) dx \equiv \mathbb{E}[x]$. To show $r \int_0^\infty e^{-rt} u(t) dt \rightarrow \lim_{t \rightarrow \infty} u(t)$ for any smooth function u , change variables $\tau = rt$: $r \int_0^\infty e^{-rt} u(t) dt = \int_0^\infty e^{-\tau} u(\tau/r) d\tau$. $u(\tau/r) \rightarrow u(\infty)$ for all $\tau > 0$. We conclude the proof by dominated convergence.

Expected earnings decomposition. Let I denote expected earnings of a given skill group. Omit skill subscripts for simplicity. Standard accounting ensures that

$$I = e^G w^G + e^C w^C + ub$$

where e^G, e^C denote the employment rates of goods producers and contractors. w^G, w^C denote the employment-weighted average wages paid by goods producers and contractors. Denote $\bar{w} = \frac{e^G w^G + e^C w^C}{e^G + e^C}$ the average wage in the economy, and $\bar{e} = e^G + e^C = 1 - u$ the employment rate. For any outcome X , we denote by X its value in a baseline equilibrium, and X' its value in the counterfactual equilibrium. Denote $\Delta X = X' - X$. Then the change in earnings between two equilibria of the model is

$$\Delta I = \underbrace{(\bar{w} - b) \Delta \bar{e}}_{\text{Employment}} + \underbrace{\left(\frac{(e^G)'}{\bar{e}'} \Delta w^G + \frac{(e^C)'}{\bar{e}'} \Delta w^C \right)}_{\text{Wage}} + \underbrace{(w^C - w^G) \Delta \left(\frac{e^C}{\bar{e}} \right)}_{\text{Wage penalty}}. \quad (45)$$

B Reduced-form results

B.1 Selection into outsourcing

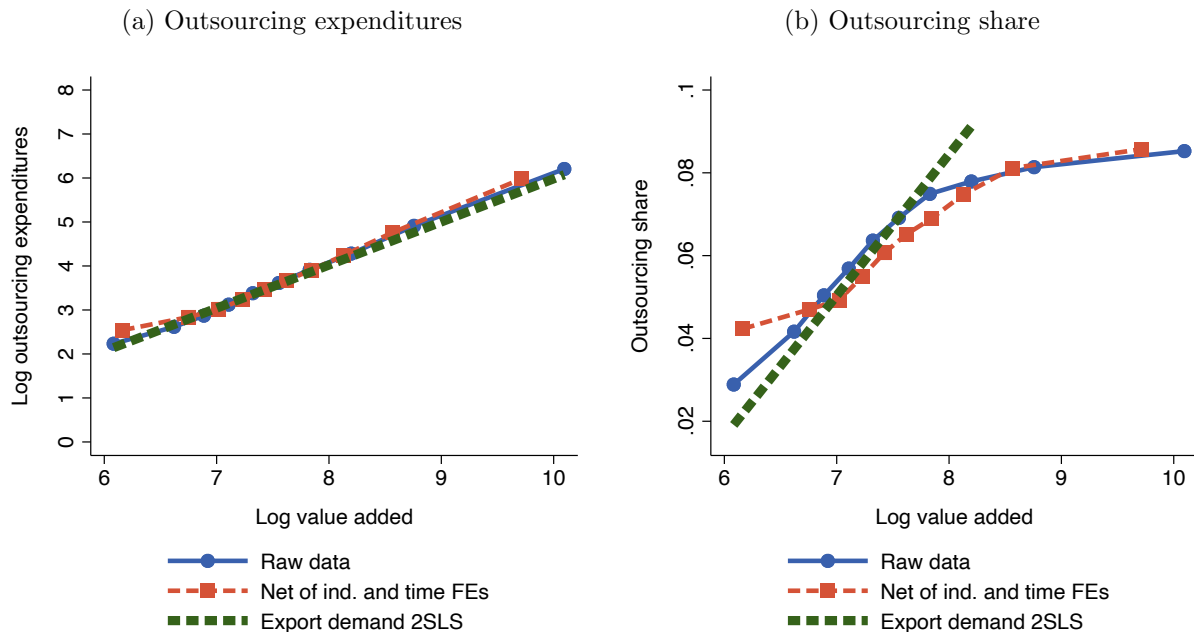
Table 4: Summary statistics

<i>Panel A: Outsourcers vs. non-outsourcers</i>				
	Non-outsourcers		Outsourcers	
	Mean	Median	Mean	Median
Employment (full-time eq.)	65	32	180	50
Sales (k)	11,133	3,543	42,149	8,189
Value added (k€)	2,979	1,270	11,696	2,454
Observations	91,608		81,939	

<i>Panel B: Outsourcing by industry</i>		
Industry	Rank	Outsourcing share
Business supplies & equipment trade	1	0.88
Telecommunications	2	0.49
Terracotta manufacturing	607	0.00
Transport into space	608	0.00

Data aggregated to three periods 1997-1999, 2000-2002, 2003-2007. Sample restricted to firms with at least 10 employees. Panel A: an outsourcer in a period is a firm that has positive expenditures on external workers in all years within the period.

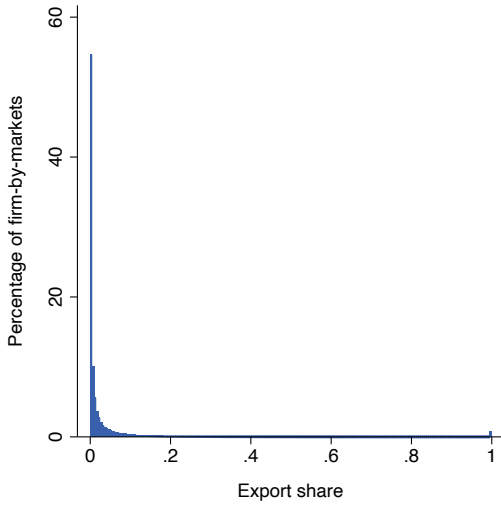
Figure 15: Outsourcing by value added.



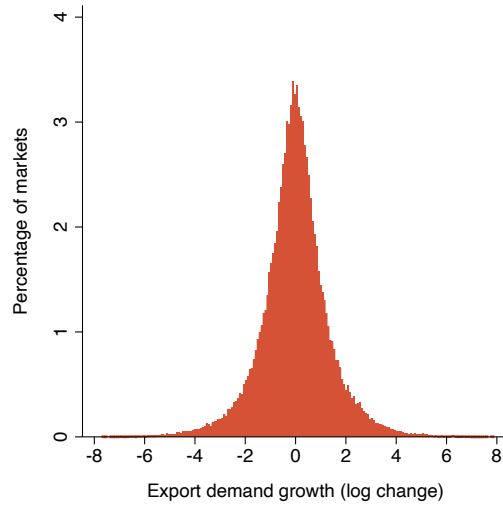
Note: Solid blue line : raw data. Dashed orange line: after removing 3-digit industry and time period fixed effects from the outsourcing share and log value added. Green line: 2SLS estimate using the export demand shift-share instrument in equation (14). Panel (a): log outsourcing expenditures. Panel (b): outsourcing share.

Figure 16: Components of export demand instrumental variable.

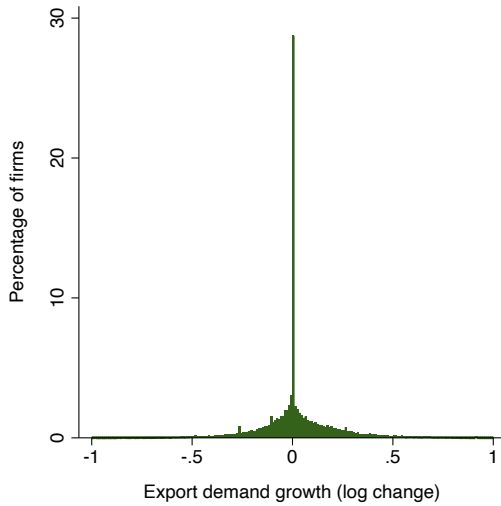
(a) Distribution of export shares by firm-market.



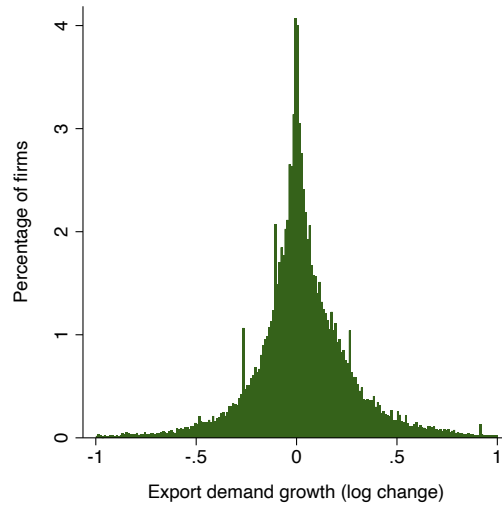
(b) Distribution of export demand by market.



(c) Distribution of IV



(d) Distribution of IV without zeros



Note: Distributions of components of the instrumental variable Z_{ft} defined in equation (14). Panel (a) shows the distribution of the firm-by-market export shares $\pi_{f,t_0,j}$. Panel (b) shows that distribution of changes in export demand $\Delta \log X_{j,t,-f}$. Panel (c) shows the distribution of the IV Z_{ft} . Panel (d) shows the same distribution after removing zeros.

Table 5: Dependent variable: log spending on external workers.

	All			Exporters				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Log V.A. + Out.	1.07*** (0.01)	1.11*** (0.01)	0.83*** (0.02)	0.81*** (0.02)	0.99*** (0.21)			
Log V.A.						0.98*** (0.22)		
Log Size							1.66*** (0.34)	
Log Labor Prod.								2.41*** (0.67)
<i>Fixed Effects</i>								
Year	✓	✓	✓	✓	✓	✓	✓	✓
Industry		✓						
Firm			✓	✓	✓	✓	✓	✓
Obs.	131734	131727	125356	38968	38968	38935	39272	38935
Stand. coef.	0.13	0.13	0.10	0.10	0.12	0.12	0.20	0.29
1 st -stage F-stat.	283.81	281.97	176.08	88.65

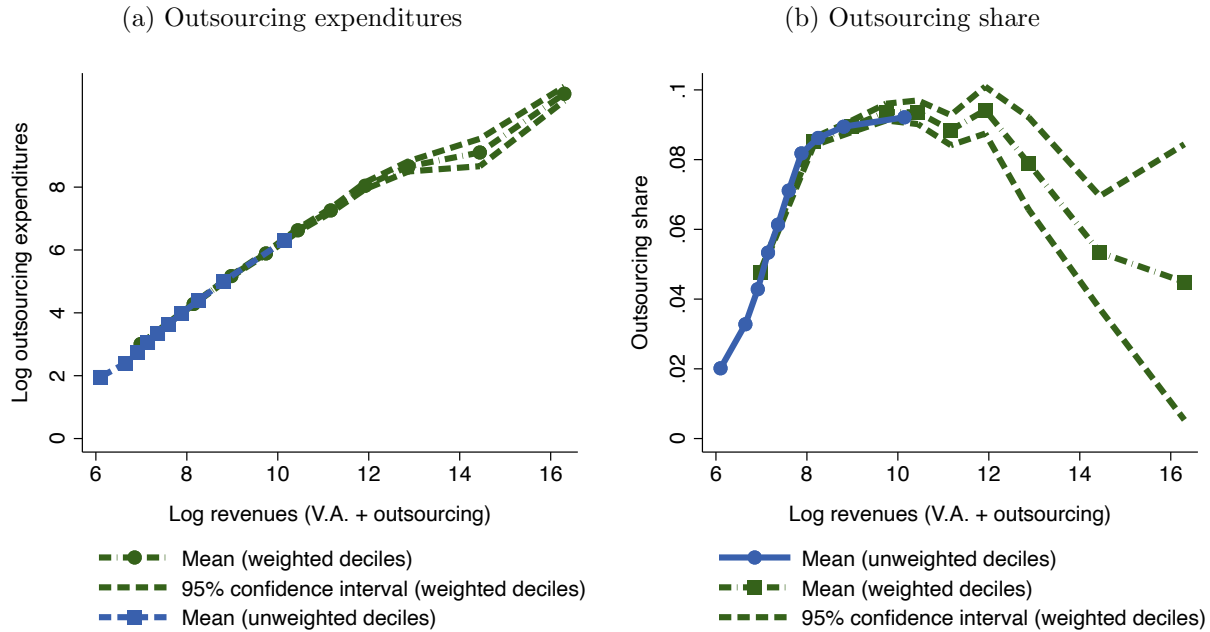
Standard errors in parenthesis, clustered by firm. ⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Dependent variable: log spending on external workers. First independent variable: log of the sum of value added and expenditures on external workers. Instrument: shift-share of export demand growth by 4-digit industry, projected by firm using firm-level export shares in first period. All regressions at firm-period level and unweighted.

Table 6: Dependent variable: spending on external workers as a fraction of labor costs, in p.p.

	All			Exporters				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Log V.A. + Out.	1.73*** (0.03)	1.82*** (0.03)	2.63*** (0.08)	1.85*** (0.12)	3.34** (1.07)			
Log V.A.						3.41** (1.09)		
Log Size							6.04*** (1.82)	
Log Labor Prod.								8.48** (2.98)
<i>Fixed Effects</i>								
Year	✓	✓	✓	✓	✓	✓	✓	✓
Industry		✓						
Firm			✓	✓	✓	✓	✓	✓
Obs.	172490	172483	172350	45798	45798	45766	46152	45766
Stand. coef.	0.21	0.22	0.32	0.22	0.40	0.41	0.73	1.03
1 st -stage F-stat.	289.56	287.65	185.67	83.57

Standard errors in parenthesis, clustered by firm. ⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Dependent variable: spending on external workers as a fraction of labor costs, in p.p. First independent variable: log of the sum of value added and expenditures on external workers. Instrument: shift-share of export demand growth by 4-digit industry, projected by firm using firm-level export shares in first period. All regressions at firm-period level and unweighted.

Figure 17: Outsourcing by revenues.



Note: Solid blue line : deciles based on number of firms. Dotted green line: deciles based on revenue-weighted quantiles. Dashed green lines: 95% confidence interval for revenue-weighted decile means. Panel (a): log outsourcing expenditures. Panel (b): outsourcing share.

B.2 The productivity effect

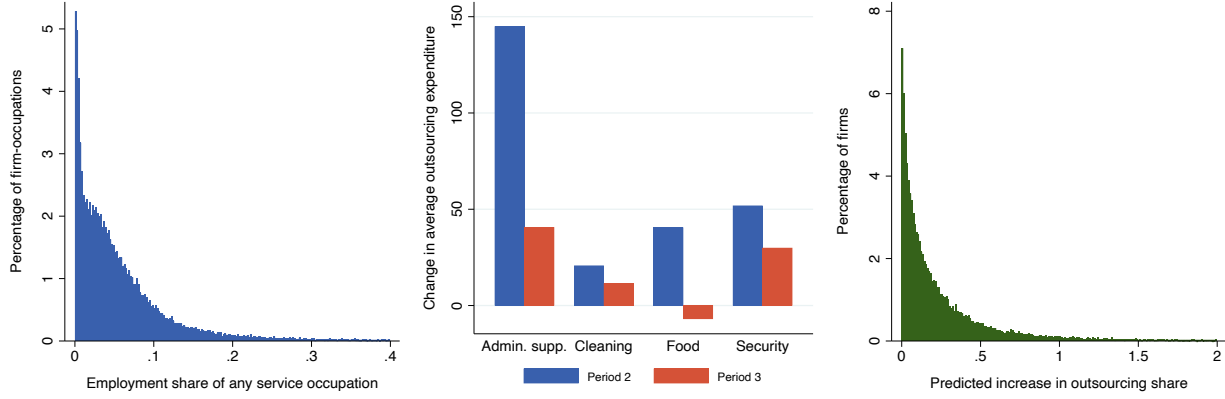
Table 7: The productivity effect of outsourcing.

	First stage	Reduced form	2SLS
	Δ out. share	Δ log VA	Δ log VA
Change in out. IV	0.245*** (0.050)	0.019*** (0.003)	
Change in out. share			0.079*** (0.020)
Obs.	46677	46677	46677
1 st -stage F-stat.			23.757

Standard errors in parenthesis, clustered by firm. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Variables winsorized at 5% level. Changes between only two periods. Instrument: shift-share of outsourcing expenditures growth by occupation, projected by firm using firm-level occupation shares in first period. Occupation codes are different in the first period and so only the second and third period can be used. Regression run in changes, leading to the lower number of observations relative to Table 6.

Figure 18: Components of outsourcing share instrumental variable.

(a) Payroll shares. (b) Change in expenditures. (c) IV



Note: Distributions of components of the instrumental variable Z'_{ft} for the firm-level outsourcing share. Panel (a) shows the distribution of the firm-by-occupation wage shares $\omega_{f,t_0,o}$ for service occupations. Panel (b) shows changes in average outsourcing expenditures by occupation $\Delta\Omega_{o,t,-f}$. Panel (c) shows the distribution of the IV Z'_{ft} . Support for Panel (c) is restricted to positive values for graphical purposes, the fraction with negative values being negligible.

B.3 The distributional effect

Table 8: Firm size wage premium in France.

Log Firm In-house Employment	0.032***	0.021***		
	(0.004)	(0.005)		
Log Firm Value Added			0.023***	
			(0.004)	
Log Firm Mean Wage				0.036*
				(0.018)
Year & 3-digit Industry Fixed Effects	✓	✓	✓	✓
Worker Fixed Effects		✓	✓	✓
Obs.	96697	94316	94316	94316

Note: Dependent variable: log worker daily wage. Standard errors in parenthesis, clustered by 3-digit industry. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Regression for service workers only, defined as in Section 3.2. In-house firm employment, value added and mean wage computed from firm-level data. Regression equation: $\log w_{i,t} = \varphi_i + \psi_{I(i,t)} + \beta X_{J(i,t)} + \eta_{i,t}$. i indexes workers, t indexes year-quarters. φ_i is a worker fixed effect. $\psi_{I(i,t)}$ is a fixed effect for the workers' employer's 3-digit industry $I(i,t)$. $J(i,t)$ denotes the worker's employer. X denotes either log employment, log value added or log mean wage.

B.4 Alternative explanations

Figure 19: Outsourcing by industry volatility.

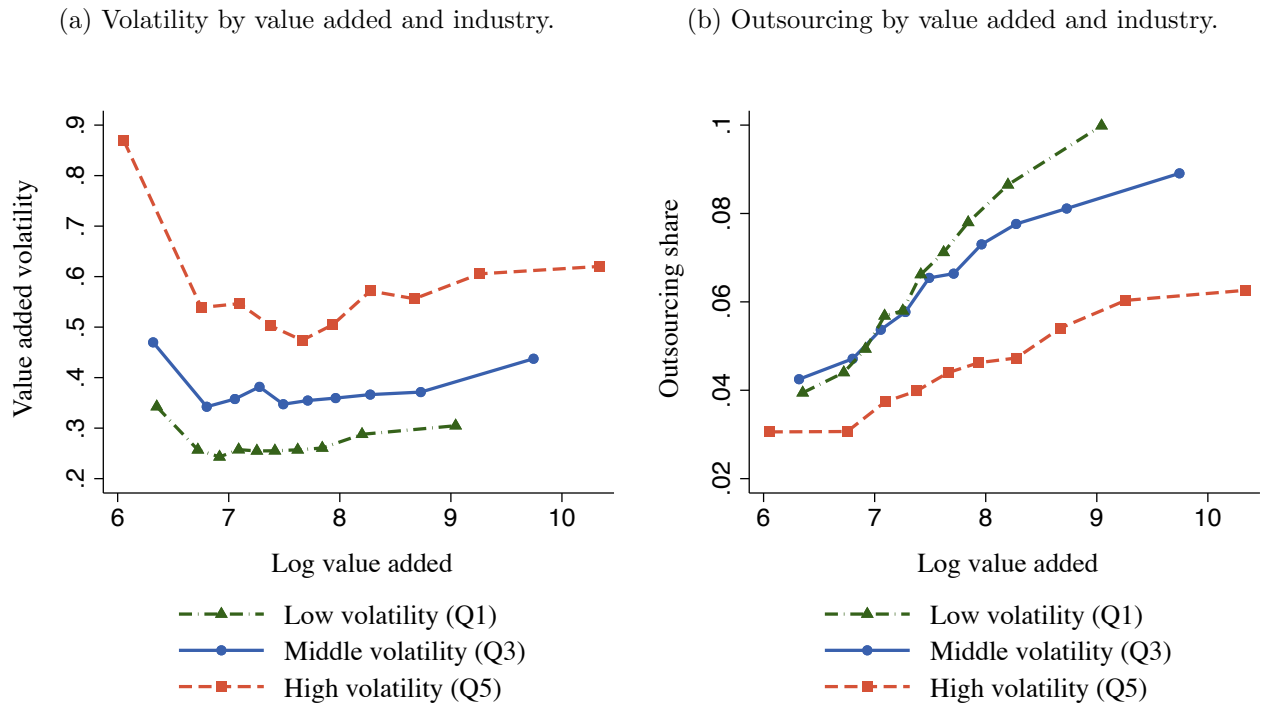


Figure 20: Size-dependent policies and outsourcing.

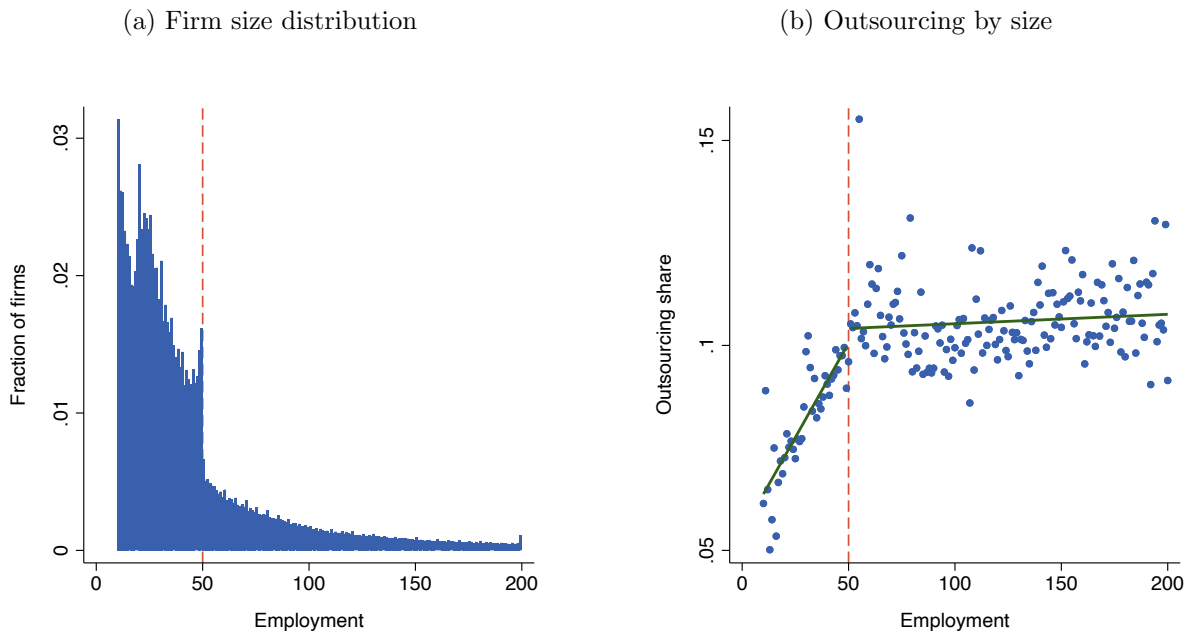


Table 9: Dependent variable: spending on external workers as a fraction of labor costs, in p.p.

	(1)	(2)	(3)	(4)	(5)	(6)
Log V.A.	1.991*** (0.017)		1.748*** (0.111)		1.858*** (0.023)	1.719*** (0.105)
Volatility		0.605*** (0.022)		0.634*** (0.036)	0.624*** (0.022)	0.707*** (0.037)
<i>Fixed Effects</i>						
Constant	✓	✓			✓	
4-digit industry			✓	✓		✓
Obs.	289243	183147	289243	183147	183147	183147
R ²	0.043	0.004	0.168	0.184	0.040	0.207

Note: Standard errors in parenthesis, clustered at the industry level. ⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Independent variables standardized to unit standard deviation. One observation is a firm.

Table 10: Dependent variable: spending on external workers as a fraction of labor costs, in p.p.

	(1)	(2)	(3)	(4)	(5)	(6)
Wage st. dev.	-0.06*** (0.00)	-0.06*** (0.00)	-0.01** (0.00)			
Wage P90-P10				-0.02*** (0.00)	-0.02*** (0.00)	-0.00* (0.00)
Log Revenues		0.02*** (0.00)	0.03*** (0.00)		0.02*** (0.00)	0.03*** (0.00)
<i>Fixed Effects</i>						
Year	✓	✓	✓	✓	✓	✓
Firm			✓			✓
Obs.	114656	113942	112716	114678	113962	112750

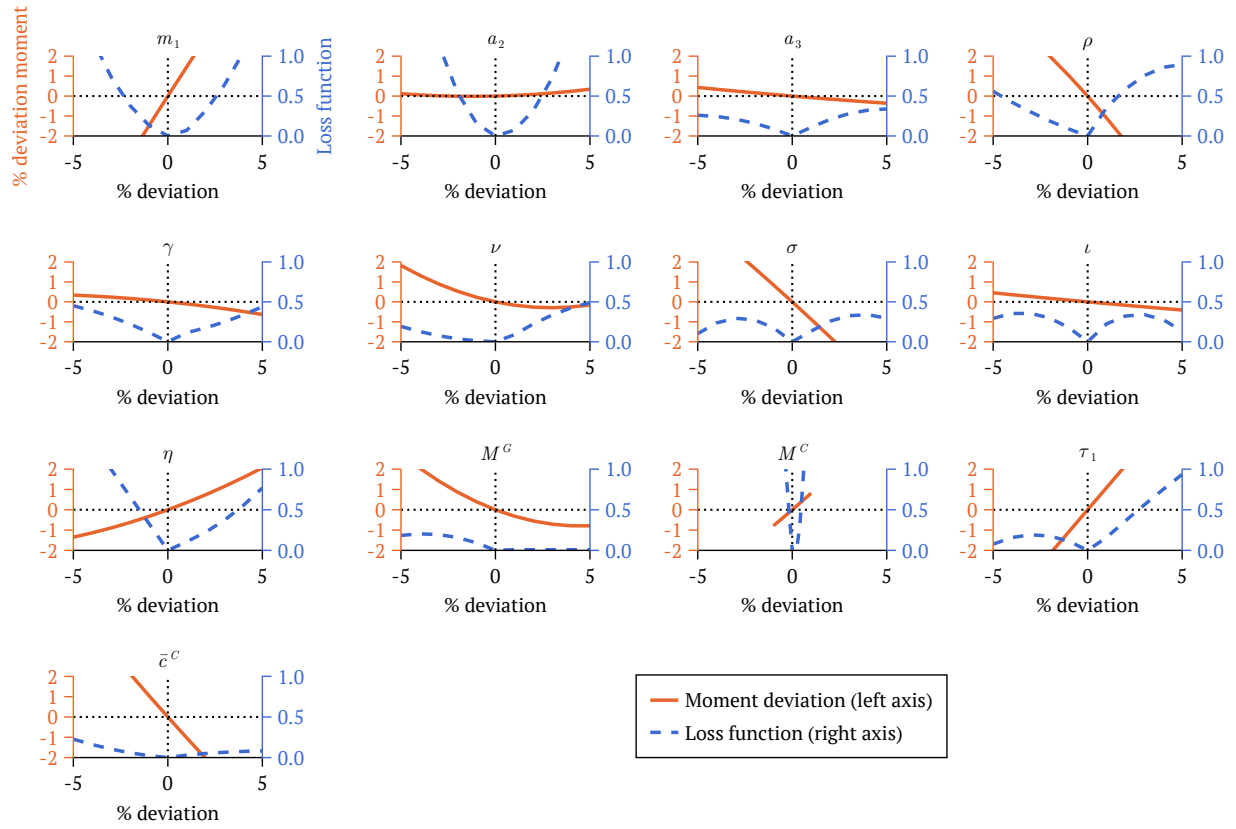
Standard errors in parenthesis, clustered by firm. ⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Dependent variable: outsourcing share, defined as expenditures on external workers divided by total expenditures on labor (gross wage bill plus expenditures on external workers). First independent variable: within-firm standard deviation of log daily wages. Second independent variable: within-firm difference between 90th percentile of log daily wages and 10th percentile of log daily wages. Revenues defined as value added plus outsourcing expenditures. All regressions at firm-period level and unweighted.

C Estimation and identification

Table 11: Parameter estimates and empirical targets.

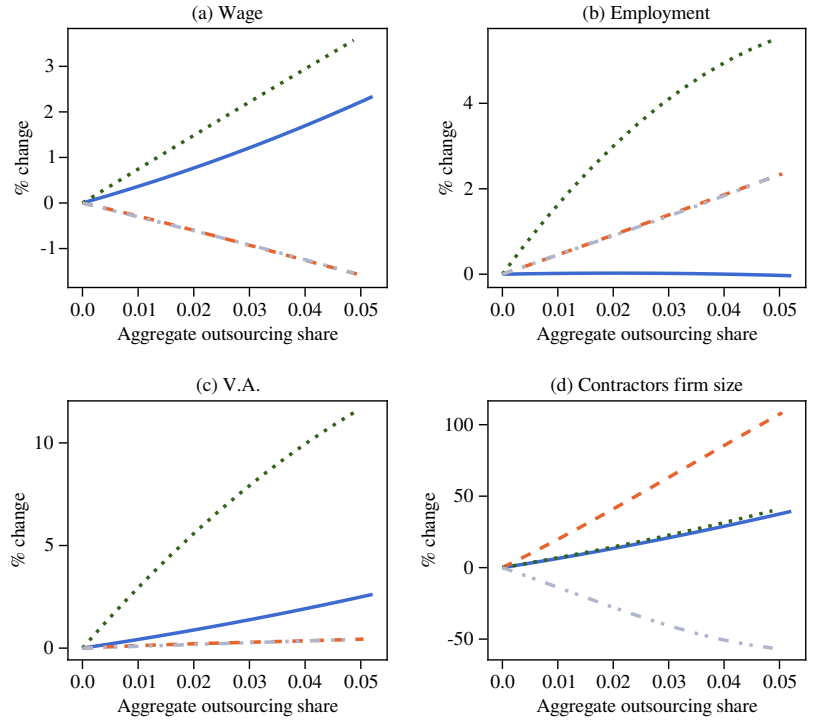
Parameter	Interpretation	Target	Empirical Moment	Simulated Moment	Parameter Estimate
A. Parameters from direct inversion.					
δ_1, δ_2	Job loss rate low-skill	EN rate low-skill	0.04		0.04
δ_3	Job loss rate high-skill	EN rate high-skill	0.03		0.03
ζ_1, ζ_2	Rel. search. emp. low-skill	EE rate low-skill	0.04		0.75
ζ_3	Rel. search. emp. high-skill	EE rate high-skill	0.03		0.68
b_1, b_2, b_3	Unemp. benefits	Replacement rate	0.70		{0.53, 0.52, 0.94}
μ_1	Matching eff. low-skill 1	NE rate low-skill	0.17		0.79
μ_2	Matching eff. low-skill 2	NE rate low-skill	0.17		0.56
μ_1	Matching eff. high-skill	NE rate low-skill	0.17		1.13
B. Parameters from MSM estimator.					
m_1	Mass of low skill 1	Out. share (employment)	0.26	0.30	0.37
a_2	Rel. prod. low-skill 2	Relative low-skill wage	1.00	1.02	0.26
a_3	Rel. prod. high-skill	Skill premium	1.74	1.78	2.62
ρ	Curvature in revenue	Labor share	0.70	0.70	0.92
γ	Curvature vac. cost	St dev. log firm size	0.98	1.03	10.1
ν	Standard dev. prod. z	St. dev. log VA	1.17	1.26	0.63
σ	Standard dev. out. costs ε	Frac. of firms OS $\leq 10\%$	0.89	0.96	0.19
ι	Correlation btw. z and ε	VA elasticity of out. share	1.83	1.63	0.15
η	Fixed cost	Frac. of firms ≤ 1 employee	0.30	0.29	0.00
M^G	Mass of producers	Average firm size	60.0	45.4	0.06
M^C	Mass of contractors	Relative contractore firm size	2.35	2.22	0.00
τ_1	Outsourcing cost	Out. share (spending)	0.08	0.08	0.65
\bar{c}^C	Providers' rel. vacancy cost	Outsourcing wage penalty	0.14	0.13	0.55

Figure 21: Simulated moments and loss function across parameter values.



Note: Numerical local identification of parameters. Solid orange line: percentage deviation of targeted moment relative to moment at estimated parameters, as a function of percentage deviation of parameter. Mapping as per Table 11. Dashed blue line: percentage deviation of loss function relative to loss function at estimated parameters, as a function of percentage deviation of parameter.

Figure 22: Effect of demand and supply shocks on equilibrium outcomes in the model.



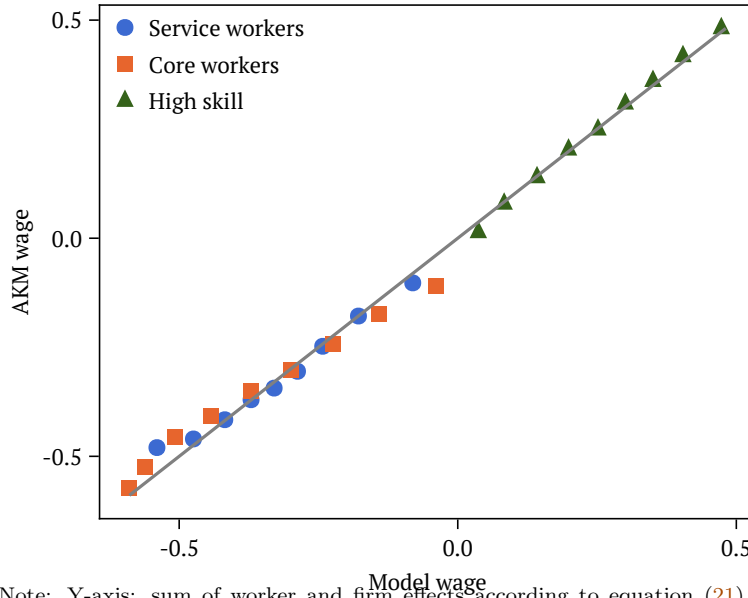
— Contractor productivity — Recruiting cost ··· Correlation - - Measure contractors
 Note: Changes in wages, employment, value added and contractor size as a function of shocks to contractor productivity τ_1 , contractor recruiting efficiency \bar{c}^C , the productivity-outsourcing cost correlation ι and the measure of contractors M^C . All changes reported as a function of the resulting increase in aggregate outsourcing share.

Table 12: Estimates of outsourcing shocks

Interpretation	Parameter	1997	2007
Contractor rel. productivity	τ_1	0.780	0.540
Correlation (z, ε)	ι	0.170	0.113
Measure of contractors	M^C	0.002	0.005
Contractor rel. vacancy cost	\bar{c}^C	1.030	0.310

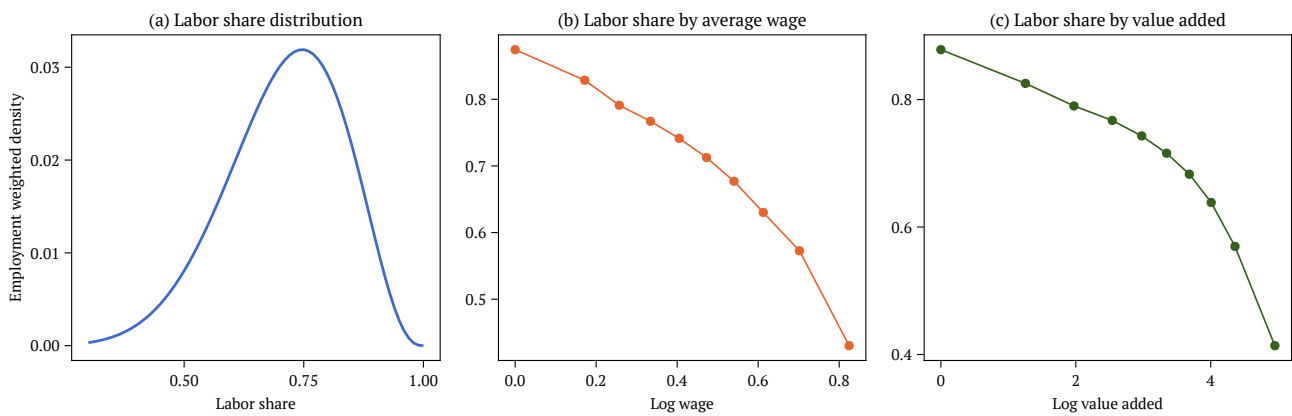
D Additional results

Figure 23: Fit between estimated worker and firm effects and wages in model.



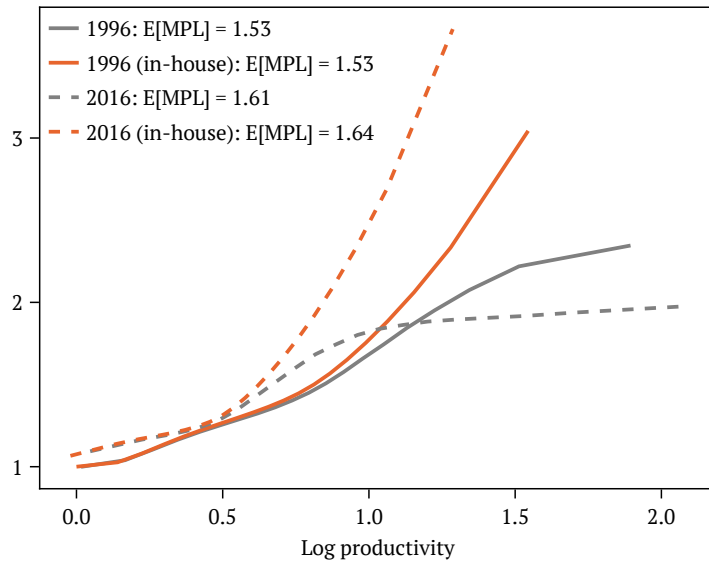
Note: Y-axis: sum of worker and firm effects according to equation (21) in model. X-axis: wages in model. Gray line: 45 degree line.

Figure 24: Rent-sharing in the labor market for all firms.




Note: Panel (a): value added-weighted distribution of labor shares. Panel (b): Labor share by average wage. Panel (c): Labor share by value added. All firms. Contractors contribute the additional mass at about Labor share = 0.9.

Figure 25: The impact of outsourcing on the marginal product of labor with all firms.



Note: Marginal product of labor for service workers. Orange lines: goods producers who hire in-house only. Gray lines: all firms, including contractors.

Figure 26: The impact of outsourcing on rent-sharing for service workers for all firms.



Note: Panel (a): decomposition of mean log wage for service workers for all firms according to the identity (24). Panel (b): distribution of the marginal product of labor for service workers for all firms. Panel (c): distribution of markdowns for service workers for all firms.

ONLINE APPENDIX

Outsourcing, Inequality and Aggregate Output

Adrien Bilal & Hugo Lhuillier

E Data description

Firm-level balance sheet data. We use the FICUS data (*“Fichier Complet Unifié de Suse”*) which covers the near universe of nonfarm French businesses. The unit of observation is a firm-year, and firms are identified by their tax identifier (*“siren”*). It details balance sheet information. We construct value added by subtracting purchases of intermediate goods and other intermediate purchases from firm sales.

Firm-level survey data. We use the EAE data (*“Enquête Annuelle d’Entreprise”*). It covers a random sample of firms and tracks them across years. We link it to other sources using the common tax identifier (*“siren”*). The unit of observation is a firm-year. Among others, the dataset breaks down intermediate purchases of goods and services. In particular, we use expenditures on external workers (*“Dépenses de personnel extérieur”*) as our main measure of outsourcing expenditures.

DADS panel. We use the 4% sample of the DADS panel, between 1996 and 2007. Once a worker enters the dataset in any year after 1976, all her subsequent employment spells are recorded. Individuals’ employment history is recorded in the dataset if (a) they have at least one employment spell, and (b) they are born in October in even years. The dataset provides start and end days of each employment spell, the job’s wage, four-digit occupation and industry, as well as establishment and firm tax identifiers that can be linked to other datasets. We follow Bilal (2021) to set sample restrictions and define unemployment.

DADS cross-section. The DADS *Postes*, are used by the French statistical institute to construct the DADS *Panel*. They cover the universe of French workers, but in the version available to researchers, worker identifiers are reshuffled every two years. The DADS *Postes* allow to compute employment, wages, occupational mix for the near universe of French establishments.

Firm-level customs data. We use customs data for the universe of French importers and exporters. The unit of observation is at the firm-product-year-country-export/import level. We aggregate French exports for every firm, year and destination country at the 4-digit industry level to construct our firm-level instrument.

F Additional proofs

F.1 Dynamic firm problem

We first show that the size constraint in (3) is consistent with the firm-level decision. Omit s indices whenever unambiguous. Denote by q the vacancy contact rate. Without loss of generality, we use a continuous offer distribution $F(w)$ to lighten notation. Start from the firm-level Kolmogorov Forward Equation:

$$\frac{dn(w, t)}{dt} = q[\phi + (1 - \phi)G(w)] - [\delta + \lambda^E(1 - F(w))]n,$$

where $\phi = \frac{u}{u + \frac{\lambda^E}{\lambda^U}(1-u)} = \frac{1}{1+k}$ is the probability of meeting an unemployed worker. In steady-state $dn/dt = 0$. Hence, from (1), $\phi + (1 - \phi)G(w) = \frac{1}{1+k(1-F(w))}$, and so $n(w) = \frac{q}{\delta} \frac{1}{[1+k(1-F(w))]^2}$. Then, from a constant returns matching function, $\lambda^U = \theta q(\theta) = \frac{M}{m[u+(1-u)\lambda^E/\lambda^U]} q(\theta)$ where θ is labor market tightness. Re-arranging leads to $q = \frac{e\delta(1+k)}{M}$. Therefore,

$$n(w) = \frac{1}{M} \frac{(1+k)e}{[1+k(1-F(w))]^2}.$$

We now turn to showing that the decisions from the dynamic profit-maximization problem of the firm coincides with those from the static firm profit maximization problem (3) when the discount rate is low enough.

Consider the dynamic problem of a firm which may be out of its long-run size, while the rest of the economy is in steady-state. Assume that firms may freely adjust their wage each instant, but face an equal-pay constraint within worker type. Without loss of generality, we consider a single worker type to make notation lighter. Firms solve

$$rJ(z, n) = \max_w R(z, n) - wn + [q(\phi + (1 - \phi)G(w)) - n(\delta + \lambda^E(1 - F(w)))]J_n(z, n).$$

Using $\phi = \frac{1}{1+k}$,

$$rJ(z, n) = \max_w R(z, n) - wn + \delta(1 + k(1 - F(w)))(n(w) - n)J_n(z, n).$$

The first-order condition implies $-n + \delta(1 + k(1 - F))n'(w)J_n + kF'(n(w) - n)J_n = 0$. Evaluated at long-run size $n = n(w)$,

$$n(w) = \delta(1 + k(1 - F))n'(w)J_n(z, n(w)).$$

The envelope condition then yields $rJ_n = R_n - w + \delta(1 + k(1 - F))[-J_n + (n(w) - n)J_{nn}]$ which again evaluated at long-run size $n = n(w)$ leads to

$$rJ_n(z, n(w)) = R_n(z, n(w)) - w - \delta(1 + k(1 - F(w)))J_n(z, n(w)).$$

When the discount rate goes to zero $r \rightarrow 0$,

$$J_n(z, n(w)) = \frac{R_n(z, n(w)) - w}{\delta(1 + k(1 - F(w)))}.$$

Substituting into the first-order condition, we obtain

$$n(w) = n'(w)(R_n - w),$$

which coincides with the static first-order condition.

F.2 Welfare and rent-sharing implications of compensating differentials models of wage inequality

Models of wage inequality based on compensating differentials have become popular due to their tractability (see e.g. Card et al., 2018, Sorkin, 2018, Berger et al., 2022, Lamadon et al., 2022). These models are well-suited to study the effect of inequality on welfare at the market level. By contrast, they require strategic interactions to generate markdown variation across firms within markets. They also struggle to translate observed wage differences into welfare difference. We make that point by considering the simplest of such models.

There are M firms j . Each firm has some productivity z_j . Each firm choose the wage w_j it posts. For simplicity, suppose that each firm is small enough that it does not consider its impact on any aggregate. Workers choose from the continuum of wage offers. Worker draw random non-wage amenities a_j specific to each firm. Workers value wages and non-work amenities a_j to maximize $V = \max_j w_j a_j$.

Tractability is achieved by assuming that a_j are independent across firms j and have a Frechet distribution with shifter T_j and shape parameter ε : $F_j(a) = e^{-T_j a^{-\varepsilon}}$.³⁷ Classical results in extreme value theory ensure several results (see e.g. Eaton and Kortum, 2002) that we derive here for completeness.

³⁷More involved correlation structures can be introduced with a nested Frechet distribution, but this additional complexity is not relevant for our purposes.

Welfare implications. First, the probability of choosing firm i is $\pi_i = \frac{T_i w_i^\varepsilon}{\sum_j T_j w_j^\varepsilon}$. To see this, consider the calculation

$$\begin{aligned}
\mathbb{P}[w_i a_i \geq w_j a_j \ \forall j \neq i] &= \int_0^\infty \prod_{j \neq i} F_j(w_i a_i / w_j) dF_i(a_i) \\
&= \int_0^\infty e^{-\Phi_{-i} w_i^{-\varepsilon} a_i^{-\varepsilon}} dF_i(a_i) \\
&= \int_0^\infty (F_i(a_i))^{\Phi_{-i} w_i^{-\varepsilon} T_i^{-1}} dF_i(a_i) \\
&= \int_0^1 F^{\Phi_{-i} w_i^{-\varepsilon} T_i^{-1}} dF \\
&= \frac{1}{1 + \Phi_{-i} w_i^{-\varepsilon} T_i^{-1}} \\
&= \frac{T_i w_i^\varepsilon}{\sum_j T_j w_j^\varepsilon}
\end{aligned}$$

where we denote $\Phi_{-i} = \sum_{j \neq i} T_j w_j^\varepsilon$, and changed variables $F = F_i(a_i)$ from the third to fourth row.

This property has made this class of models particularly attractive to model an upward-sloping labor supply curve. When every firm is small enough that it neglects its effect on the denominator, the labor supply curve faced by every firm j is of the form $n_j(w) = n_0 T_j w^\varepsilon$.

This tractability comes, however, at a cost. The cost is that welfare is equalized across firms, regardless of the wage they decide to post. To see this second result (also well-established, but less well-known), compute as above

$$\begin{aligned}
\mathbb{P}[w_i a_i \leq \omega \ \& \ w_i a_i > w_j a_j \ \forall j \neq i] &= \int_0^{\omega/w_i} \prod_{j \neq i} F_j(w_i a_i / w_j) dF_i(a_i) \\
&= \int_0^{\omega/w_i} e^{-\Phi_{-i} w_i^{-\varepsilon} a_i^{-\varepsilon}} dF_i(a_i) \\
&= \int_0^{\omega/w_i} (F_i(a_i))^{\Phi_{-i} w_i^{-\varepsilon} T_i^{-1}} dF_i(a_i) \\
&= \int_0^{F_i(\omega/w_i)} F^{\Phi_{-i} w_i^{-\varepsilon} T_i^{-1}} dF \\
&= \frac{F_i(\omega/w_i)^{1 + \Phi_{-i} w_i^{-\varepsilon}}}{1 + \Phi_{-i} w_i^{-\varepsilon} T_i^{-1}} \\
&= \frac{T_i w_i^\varepsilon}{\sum_j T_j w_j^\varepsilon} e^{-\Phi \omega^{-\varepsilon}}
\end{aligned}$$

where $\Phi = \sum_j T_j w_j^\varepsilon$. Therefore, the conditional distribution of values within every chosen firm i is

$$\begin{aligned}
\mathbb{P}[w_i a_i \leq \omega \mid w_i a_i > w_j a_j \ \forall j \neq i] &= \frac{\mathbb{P}[w_i a_i \leq \omega \ \& \ w_i a_i > w_j a_j \ \forall j \neq i]}{\mathbb{P}[w_i a_i \geq w_j a_j \ \forall j \neq i]} \\
&= e^{-\Phi \omega^{-\varepsilon}}
\end{aligned}$$

The distribution of indirect utility of employed workers at any firm is Frechet with shifter Φ and shape parameter ε . It is independent of the wage offered by the firm. Thus, welfare is equalized across firms. In particular, expected utility of employed workers at any firm is $\Gamma(1 - 1/\varepsilon)\Phi^{1/\varepsilon}$, where Γ denotes Euler's Gamma function. Non-work amenities introduce mixing, but they do not change the core implications relative to a standard free-mobility condition without amenities, $\max_j w_j$.

Markdowns and labor shares. To study rent-sharing in this economy, we specify a revenue function $R(z, n)$ that may depend on productivity z and the number of workers that a firm hires n . Any given firm j then solves

$$\max_{w,n} R(z_j, n) - wn \text{ s.t } n \leq \pi_j(w) \equiv \frac{T_j w^\varepsilon}{\sum_{i \neq j} T_i w_i^\varepsilon}$$

Assuming a large enough number of firms and thus no strategic interactions, $\sum_{i \neq j} T_i w_i^\varepsilon \equiv W_{0j}^\varepsilon T_j$ may be assumed to be constant from the perspective of firm j . The labor supply curve becomes $\pi_j(w) = W_{0j}^{-\varepsilon} w^\varepsilon$. The wage it pays is $w_j(n) = W_{0j} n^{1/\varepsilon}$. The problem of the firm becomes

$$\max_n R(z_j, n) - W_{0j} n^{1+1/\varepsilon}.$$

Thus, the firm chooses n_j such that $R_n(z_j, n) = (1 + 1/\varepsilon)W_{0j} n_j^{1/\varepsilon}$. Under isoelastic revenue $R(z, n) = zn^\rho$, $R_n(z_j, n_j) = \rho R(z_j, n_j)/n_j = \rho z_j n_j^{\rho-1} = (1 + 1/\varepsilon)W_{0j} n_j^{1/\varepsilon} = (1 + 1/\varepsilon)w_j(n_j)$.

Hence, the markdown is constant across firms and given by

$$\text{markdown} = \frac{\rho\varepsilon}{1 + \varepsilon}.$$

The labor share is also constant across firms and given by

$$\text{LS} = \frac{\varepsilon}{1 + \varepsilon}$$

Taking stock. Thus, it is difficult to connect models of wage inequality based on compensating differentials to welfare inequality across firms.³⁸ These results are specific to the Frechet distribution however. Potentially, more complex distribution of compensating differentials could lead to meaningful welfare effects. Yet, more complex distributions would also break the tractability provided by the Frechet distribution. Similarly, models of wage inequality based on compensating differentials deliver constant markdowns and labor shares under standard isoelastic revenue functions absent market power.

³⁸Of course, the shifter Φ depends on the distribution of wages inside the market and so cross-market welfare inequality is still meaningful.

G Simulation and estimation

G.1 Model

Good producers Good producers solve the following problem

$$\pi(z, \varepsilon) = \max_{\mathbf{n}, \mathbf{w}, \mathbf{v}, o} R(z, \mathbf{n}) - ((1 - o)w_1 - op\varepsilon)n_1 - (1 - o)c(v_1) - \sum_{s>1} (w_s n_s + c(v_s)) - \eta, \quad (46)$$

subject to

$$n_s \leq n_s(w) \left(\frac{v_s}{V_s} \right) \quad \text{if } o_s = 0,$$

where $\mathbf{x} \equiv \{x_s\}_{s=1}^S$ and n_s is the firm-specific labor supply curve,

$$n_s(w) \equiv \frac{(1 + k_s)e_s}{(1 + k_s(1 - F_s(w_s)))^2}.$$

In the formulation of the problem above, we have already assumed that only the workers of skill type one (service workers) can be outsourced. We use the following parametric assumptions. The revenue function is Cobb-Douglas nested in a decreasing returns upper tier,

$$R(z, \mathbf{n}) = \left(z \prod_{s=1}^S n_s^{a_s} \right)^\rho$$

with $\sum_s a_s = 1$. The vacancy cost function is isoelastic with elasticity γ ,

$$c(v) = \frac{v^{1+\gamma}}{1 + \gamma},$$

where we have normalized the vacancy cost of a vacancy for good producers to one, $c_0 = 1$. Finally, the joint distribution of (z, ε) , denoted by Γ , is log-normal with zero means and variance-covariance matrix Σ .

Contractors To solve numerically the contractors' problem in a similar fashion as the problem of good producers, we introduce some minimal productivity heterogeneity across contractors. Specifically, let Ψ denote the distribution of contractors' productivity. We parametrize it to be log-normal with zero mean and variance σ^C and we set $\sigma^C \approx 0$. Then, contractors solve the problem

$$\pi^C(z) = \max_{n, w, v} p\tau zn - wn - \frac{(\bar{c}^C)^\gamma v^{1+\gamma}}{1 + \gamma} \quad \text{s.t.} \quad n \leq n_1(w) \left(\frac{v}{V_1} \right), \quad (47)$$

where z denote the productivity of the contractor. The effective amount of labor that a contractor with productivity z provides to the outsourcing market is τzn . Let \underline{z}^C be the least productive contractor

firms. In equilibrium, this firm offers the reservation wage and $\underline{z}^C = \underline{w}_1/(p\tau)$.³⁹

G.2 A tractable reformulation of the model

Due to its large amount of heterogeneity, this model is *a priori* numerically non-tractable. In particular, two objects are complicated to compute. First, the wage offered to in-house workers by good producers that outsource their service workers. These wages depend indeed on (z, ε) rather than on a uni-dimensional variable. Second, the wage offer distributions is non-trivial to compute since different types of firms are now competing on the same skill-specific job ladder.⁴⁰ In this section, we derive a reformulation of the problem that simplifies these two problems. This derivation is feasible under three assumptions: a Cobb-Douglas revenue function, a single outsourceable worker type, and a log-normal distribution for (z, ε) .

Outsourcing good producers Index good producers that outsource their service workers by the superscript o . Similarly, index good producers that hire their service workers in-house by i . The problem of a good producer of type o reads

$$\pi^o(z, \varepsilon) = \max_{\mathbf{n}, \{w_s\}_{s>1}, \{v_s\}_{s>1}} R(z, \mathbf{n}) - n_1 p \varepsilon - \sum_{s>1} (w_s n_s + c(v_s)) - \eta. \quad (48)$$

The optimal number of service workers hired from the contractor sector is obtained by taking the first order condition of this problem with respect to n_1 ,

$$n_1^o(z, \varepsilon) = \left(\frac{\rho a_1}{p \varepsilon} \right)^{\frac{1}{1-a_1 \rho}} \left(z \prod_{s>1} n_s^{a_s} \right)^{\frac{\rho}{1-a_1 \rho}}.$$

Plugged back into (48), the profit of the firm rewrites

$$\pi^o(z, \varepsilon) = \mathcal{G} \left(z \varepsilon^{-a_1} \prod_{s>1} n_s^{a_s} \right)^\kappa - \sum_{s>1} w_s n_s - c(v_s) - \eta \equiv \pi^o(\hat{z}),$$

for κ a parametric constant and \mathcal{G} a general equilibrium constant.⁴¹ In the above expression, payroll and vacancy costs are independent from z and ε . Meanwhile, revenues only depend on the TFP

³⁹When contractor firms post their optimal number of vacancy, their profits are given by

$$\pi^C(z) \propto \max_w [(p\tau z - w)n_1(w)]^{\frac{1+\gamma}{\gamma}}.$$

Hence, $\underline{z}^C = w(\underline{z}^C)/(p\tau)$. Furthermore, the least productive active contractor firms must offer the reservation wage for otherwise contractor firms with lower productivity would be able to make a profit by posting a wage in $[\underline{w}, w(\underline{z}^C)]$.

⁴⁰In a standard Burdett-Mortensen model, the wage offer distribution is directly recovered from two differential equations obtained from the wage and vacancy first-order conditions. This is not the case here as two firms with a similar revenue TFP z may offer different wages depending on their outsourcing choice.

⁴¹Specifically, we have

$$\kappa = \frac{\rho}{1 - \rho a_1} \quad \text{and} \quad \mathcal{G} = \frac{\rho}{\kappa} \left(\frac{\rho a_1}{p} \right)^{a_1 \kappa}.$$

aggregator $\hat{z} \equiv z\varepsilon^{a_1}$. As a result, the policy functions of type- o good producers are only a function of \hat{z} and it is not needed to keep track of z and ε separately. Since (z, ε) is jointly log-normally distributed, so is (z, \hat{z}) and a closed-form expression exists for its variance-covariance matrix. Let Φ denote the log-normal distribution under the change of variable $(z, \varepsilon) \rightarrow (z, \hat{z})$. Finally, profits of in-house and outsourcing good producers are increasing in z and \hat{z} respectively, and there exists two productivity lower bounds, \underline{z} and $\underline{\hat{z}}$, so that $\pi^i(z) > 0$ iff $z > \underline{z}$ and $\pi^o(\hat{z}) > 0$ iff $\hat{z} > \underline{\hat{z}}$. With a slight abuse of notation, we refer to Φ as the truncated log-normal distribution, Φ_z as the z -marginal, and $\Phi_{\hat{z}|z}$ as the distribution of \hat{z} condition on z .

Wage distributions To derive the wage offer distributions, it is required to know which firms outsource their service workers and which do not. Let $\varphi(z)$ denote the productivity level that renders an outsourcing firm indifferent between the two outsourcing choices, $\pi^i(z) = \pi^o[\varphi(z)]$ and $\pi^i(z) < \pi^o(\hat{z})$ for all $\hat{z} > \varphi(z)$. With this notation, the wage offer distribution of skill $s > 1$ is given by

$$F_s(w) = \frac{M^G}{V_s} \left(\int \mathbf{1}\{w_s^i(z) \leq w\} v_s^i(z) \Omega^i(z) d\Phi_z(z) + \int \mathbf{1}\{w_s^o(\hat{z}) \leq w\} v_s^o(\hat{z}) \Omega^o(\hat{z}) d\Phi_{\hat{z}|z}(\hat{z}) \right), \quad (49)$$

where

$$V_s = M^G \left(\int v_s^i(z) \Omega^i(z) d\Phi_z(z) + \int v_s^o(\hat{z}) \Omega^o(\hat{z}) d\Phi_{\hat{z}|z}(\hat{z}) \right) \quad (50)$$

is the mass of vacancy posted for skill s and $\Omega^i(x) \equiv \Phi_{z|x}[\varphi(x) | x]$ is the probability that good producer z hires its service workers in-house. Similarly, $\Omega^o(x) \equiv \Phi_{z|\hat{z}}[\varphi^{-1}(x) | x]$. The first integral in (49) is the relative mass of vacancy attached to wages lower than w and offered by good producers hiring their service workers in-house. The second integral is the relative mass of vacancy attached to wages lower than w and offered by good producers outsourcing their service workers. Similarly, the wage offer distribution for service workers is

$$F_1(w) = \frac{M^G}{V_1} \int \mathbf{1}\{w_1^i(z) \leq w\} v_1^i(z) \Omega^i(z) d\Phi_z(z) + \frac{M^C}{V_1} \int \mathbf{1}\{w_1^C(z) \leq w\} v_1^C(z) d\Psi(z), \quad (51)$$

where

$$V_1 = M^G \int v_1^i(z) \Omega^i(z) d\Phi_z(z) + M^C \int v_1^C(z) d\Psi(z). \quad (52)$$

In equation (51), the first integral is the relative mass of vacancy attached to wages lower than w and offered by good producers hiring their service workers in-house. The second integral is the mass of vacancy attached to wages lower than w and offered by contractors.

Finally, two further conditions close the equilibrium. First, the reservation wages which constitute the lower bound of the wage offer distributions,

$$\underline{w}_s = b_s + (k_s^U - k_s) \left(\int_{\underline{w}_s} \frac{1 - F_s(w)}{1 + k_s(1 - F_s(w))} dw \right). \quad (53)$$

Second, the market clearing condition of the outsourcing market is

$$\begin{aligned} M^C \int \tau z n^c(z) d\Psi(z) &= M^G \int \varepsilon n_1^o(z, \varepsilon) \mathbf{1}\{\pi^o(z, \varepsilon) > \pi^i(z)\} d\Phi(z, \varepsilon) \\ &= M^G \left(\frac{a_1 \rho}{p} \right)^{\frac{\kappa}{\rho}} \int \left(\hat{z} \prod_{s>1} n_s^o(\hat{z})^{a_s} \right)^{\kappa} \Omega^o(\hat{z}) d\Phi_{\hat{z}}(\hat{z}), \end{aligned} \quad (54)$$

where, as explained in the main text, ε is interpreted as an idiosyncratic iceberg shock and as such appears in the aggregate demand for outsourced service workers. The second line is obtained from plugging the expression for n^o and performing the change of variable $z\varepsilon^{-a_1} \rightarrow \hat{z}$.

Definition 1 (Equilibrium). *An equilibrium is a collection of wage and vacancy functions for good producers, $\{w_s^\theta, v_s^\theta\}_{s \in \{1, \dots, S\}, \theta \in \{i, o\}}$, wage and vacancy functions for contractor firms, w_1^C and v_1^C , an indifference function, φ , wage distributions $\{F_s\}_{s=1}^S$, productivity cutoffs $(\underline{z}, \hat{z}, \underline{z}^C)$, and aggregate quantities, $\{\underline{w}_s\}_{s=1}^S$ and p , such that*

1. Given $\{F_s\}_{s=1}^S$, p , and an outsourcing decisions $\theta \in \{i, o\}$, the functions $\{w_s^\theta, v_s^\theta\}_{s=1}^S$ solve (46);
2. Given F_1 and p , the functions w_1^C and v_1^C solve (47);
3. Given the policy functions $\{w_s^\theta, v_s^\theta\}_{s=1}^S$, the indifference function is such that $\pi^i(z) = \pi^o[\varphi(z)]$;
4. Given the policy functions, the wage distributions satisfy (49) and (51);
5. Given $\{F_s\}_{s=1}^S$, the reservation wages are given by (53);
6. Given the firms' profits, the productivity cutoffs solve $\underline{z} = \inf\{z : \pi^i(z) > 0\}$, $\hat{z} = \inf\{\hat{z} : \pi^o(\hat{z}) > 0\}$ and $\underline{z}^C = \underline{w}_1 / (p\tau)$;
7. Given the policy functions and φ , the price p solves the market clearing condition (54).

G.3 Expressing the model as a system of differential equations

To compute numerically the equilibrium defined in Definition 1, we take the first order conditions of the firms' problem (46) and (47). In doing so, we show that it is never required to solve for the solution of the contractors' problem as long as their wage overlap with the wages offered by good producers. We then rewrite the wage offer distributions (49) and (51) in differentials so as to obtain them jointly with the wage and vacancy policy functions as the solution to a system of differential equations.

Optimality conditions Let $\Upsilon_s^\theta(z) \equiv F_s[w_s^\theta(z)]$ denote the rank on the skill- s job ladder of a good producer with productivity z and outsourcing choice $\theta \in \{i, o\}$. For instance, for $s > 1$ and $\theta = i$, we have

$$\Upsilon_s^i(z) = \frac{M^G}{V_s} \left(\int^z v_s^i(x) \Omega^i(z) d\Phi_z(x) + \int \mathbf{1}\{w_s^o(\hat{z}) \leq w_s^i(z)\} v_s^o(\hat{z}) \Omega^o(\hat{z}) d\Phi_z(\hat{z}) \right). \quad (55)$$

Taking the first-order condition of (46) with respect to w and using the function Υ_s^θ , the wage optimality condition can be expressed as the differential equation

$$\frac{\partial w_s^\theta(z)}{\partial z} = \left(\frac{2k_s}{1 + k_s(1 - \Upsilon_s^\theta(z))} \right) \left(\frac{\partial R(z, \mathbf{n}^\theta(z))}{\partial n_s} - w_s^\theta(z) \right) \frac{\partial \Upsilon_s^\theta(z)}{\partial z}, \quad (56)$$

while the vacancy optimality condition is

$$v_s^\theta(z)^\gamma = n_s^\theta(z) \left(\frac{\partial R(z, \mathbf{n}^\theta(z))}{\partial n_s} - w_s^\theta(z) \right). \quad (57)$$

Similarly, the wage optimal condition for a contractor firm is

$$\frac{\partial w_1^C(z)}{\partial z} = \left(\frac{2k_1}{1 + k_1(1 - \Upsilon_1^C(z))} \right) (p\tau z - w_1^C(z)) \frac{\partial \Upsilon_1^C(z)}{\partial z}, \quad (58)$$

while the vacancy optimality condition is

$$(\bar{c}^C v_1^C(z))^\gamma = n_1^C(z) (p\tau z - w_1^C(z)). \quad (59)$$

Solving for the wage in (56) and (58) requires to know $d\Upsilon_s^\theta$. The goal of the algorithm developed in Section G.4 is to solve jointly for w_s^θ and Υ_s^θ . This requires to know how the second integral in (55) moves with z . To deal with this problem, we define two wage equivalence functions. First, let $\zeta^{G \rightarrow C}$ be the TFP of contractor firms posting wage $w_1^i(z)$; that is, $w_1^i(z) = w_1^C[\zeta^{G \rightarrow C}(z)]$. Similarly, let $\zeta_s^{i \rightarrow o}$ be such that $w_s^i(z) = w_s^o[\zeta_s^{i \rightarrow o}(z)]$. We now derive expressions for these two functions.

Service workers' wages Given that wages are strictly increasing in productivity, the function $\zeta^{G \rightarrow C}$ is well-defined on the joint support of the wages offered by good producers and contractor firms. Since z is unbounded above for both good producers and contractor firms, $w_1^i(z) \rightarrow \bar{w}_1$ and $w_1^C(z) \rightarrow \bar{w}_1$ as $z \rightarrow \infty$. In addition, we have already argued that the least productive contractor firm offers the reservation wage. It follows that the function $\zeta^{i \rightarrow o}$ is well-defined for $z \geq \underline{z}$. Furthermore, a closed-form expression exists for $\zeta^{i \rightarrow o}$. For all $z \geq \underline{z}$, $w_1^i(z) = w_1^C[\zeta^{G \rightarrow C}(z)]$ implies by definition $\Upsilon_1^i(z) = \Upsilon_1^C[\zeta^{G \rightarrow C}(z)]$ and therefore $n_1^i(z) = n_1^C[\zeta^{G \rightarrow C}(z)]$. Furthermore, differentiating the first two equations implies $\partial_z w_1^i(z) / \partial_z \Upsilon_1^i(z) = \partial_z w_1^C[\zeta^{G \rightarrow C}(z)] / \partial_z \Upsilon_1^C[\zeta^{G \rightarrow C}(z)]$, where $\partial_z x(z) \equiv dx/dz$. Comparing equations (56) and (58), it must then be that the marginal product of labor of good producer z and contractor firms $\zeta^{G \rightarrow C}(z)$ are equal. But since contractor firms are facing constant returns to scale, their MPL are independent from their size, and we can therefore invert this MPL

equality condition and solve for $\zeta^{G \rightarrow C}(z)$ to obtain

$$\zeta^{G \rightarrow C}(z) = \left(\frac{\rho a_1}{p\tau} \right) \left(\frac{R^i(z)}{n_1^i(z)} \right). \quad (60)$$

Finally, since the right hand side of (57) and (59) are equal, it must also be that

$$v_1^i(z) = \bar{c}^C v_1^C[\zeta^{G \rightarrow C}(z)].$$

With the help of these results, we can differentiate $\Upsilon_1^i(z)$ for $z \geq \underline{z}$ to obtain

$$\frac{\partial \Upsilon_1^i(z)}{\partial z} = \frac{v_1^i(z)}{V_1} \left[M^G \Omega^i(z) \frac{\partial \Phi_z(z)}{\partial z} + \left(\frac{M^C}{\bar{c}^C} \right) \frac{d\Psi[\zeta^{G \rightarrow C}(z)]}{dz} \right]. \quad (61)$$

Ignoring for now the complementarity across skills, equations (56) and (61) constitute a system of differential equations, which, together with (57) and (60), allows us to solve for $(w_1^i, v_1^i, \Upsilon_1^i)$. These differential equations are subject to two boundary conditions, $\underline{w}_1^i \equiv w_1^i(\underline{z})$ and $\underline{\Upsilon}_1^i \equiv \Upsilon_1^i(\underline{z})$. If the least productive good producers decide to offer the reservation wage, then $\underline{w}_1^i = \underline{w}_1$ and $\underline{\Upsilon}_1^i(\underline{z}) = 0$. However, since good producers compete on the job ladder of service workers with the contractor firms, the least productive good producers may decide optimally *not* to offer the reservation wage.⁴² If that is the case, then all wages between $[\underline{w}_1, \underline{w}_1^i)$ are offered by contractor firms. These wages have to satisfy (58) together with

$$\frac{\partial \Upsilon_1^C(z)}{\partial z} = \left(\frac{M^C}{V_1} \right) v_1^C(z) \frac{\partial \Psi(z)}{\partial z}.$$

Combined with (59), this differential equation becomes

$$\frac{\partial \Upsilon_1^C(z)}{\partial z} = \frac{1}{\bar{c}^C} ((p\tau z - w_1^C(z)) n_1^C(z))^{1/\gamma} \left(\frac{N}{V_1} \right) \frac{\partial \Psi(z)}{\partial z}. \quad (62)$$

This differential equation, together with (58) and the boundary conditions $w_1^C(\underline{z}^C) = \underline{w}_1$ and $\Upsilon_1^C(\underline{z}^C) = 0$, pin down the wages and the wage offer distribution for contractors with TFP in $[\underline{z}^C, \tilde{z}^C)$, where \tilde{z}^C is such that $w_1^C(\tilde{z}^C) = w_1^i(\underline{z})$. By a similar argument as before, this wage equality also implies that the MPL of the two firms have to be equal, or

$$\tilde{z}^C = \left(\frac{\rho a_1}{p\tau} \right) \left(\frac{R^i(\underline{z})}{n_1^i(\underline{z})} \right). \quad (63)$$

Hence, on $[\underline{z}^C, \tilde{z}^C)$, the policy functions of contractor firms are obtained from solving a standard uni-skill constant returns to scale Burdett-Mortensen model. For $z \geq \tilde{z}^C$, it is not needed to solve for the problem of contractor firms since we know that $w_1^C[\zeta^{G \rightarrow C}(z)] = w_1^i(z)$ and $v_1^C[\zeta^{G \rightarrow C}(z)] = v_1^i(z)/\bar{c}^C$.

⁴²Precisely, the least productive good producers may not want to offer the reservation wage due to two reasons. First, due to the presence of the contractor firms on the job ladder. Second, due to the existence of decreasing returns to scale in the revenue function. If either feature were absent, then $\underline{w}_1^i = \underline{w}_1$ as in the standard Burdett-Mortensen model.

Other skills While only good producers are hiring workers with skills $s > 1$, there is effectively two type of goods producers in this model: those that hire service workers in-house, and those that outsource their service workers. As such, a similar derivation as in the previous paragraph is needed to solve efficiently for the wage distribution of skills $s > 1$. To avoid complicated combinatorial issues, we assume that the least productive good producers post the reservation wage regardless of their outsourcing decision; that is, $w_s^i(z) = w_s^o(\hat{z}) = \underline{w}_s$ for $s > 1$.⁴³ With this assumption, the wage equivalence function $\zeta^{i \rightarrow o}$ is globally well-defined. Then, using a similar argument as in the previous paragraph, $w_s^i(z) = w_s^o[\zeta^{i \rightarrow o}(z)]$ implies $v_s^i(z) = v_s^o[\zeta^{i \rightarrow o}(z)]$ as well as the equalization of the MPL of these two firms. Together, we therefore know that

$$\frac{R^i(z)}{n_s^i(z)} = \left(\frac{\kappa}{\rho}\right) \left(\frac{R^o[\zeta_s^{i \rightarrow o}(z)]}{n_s^o[\zeta_s^{i \rightarrow o}(z)]}\right) = \left(\frac{\kappa}{\rho}\right) \left(\frac{R^o[\zeta_s^{i \rightarrow o}(z)]}{n_s^i(z)}\right), \quad (64)$$

where the second equality follows from $n_s^i(z) = n_s^o[\zeta_s^{i \rightarrow o}(z)]$ and $v_s^i(z) = v_s^o[\zeta_s^{i \rightarrow o}(z)]$ so that $n_s^i(z) = n_s^o[\zeta_s^{i \rightarrow o}(z)]$. Hence, it must be that $R^o[\zeta_s^{i \rightarrow o}(z)] = \rho R^i(z)/\kappa$, which holds for each $s > 1$. This in turn implies $R^o[\zeta_s^{i \rightarrow o}(z)] = R^o[\zeta_{s'}^{i \rightarrow o}(z)]$ for any pair of skill (s, s') . But the function R^o is strictly increasing in z , and therefore it must be that $\zeta_s^{i \rightarrow o}(z) = \zeta_{s'}^{i \rightarrow o}(z) \equiv \zeta^{i \rightarrow o}(z)$. The skill independence of the function $\zeta^{i \rightarrow s}$ allows us to simplify the equation (64) to

$$\zeta^{i \rightarrow o}(z) = \left(\frac{\rho}{\mathcal{G}\kappa}\right)^{\frac{1}{\kappa}} \left(z n_1^i(z)^{a_1} \prod_{s>1} n_s^i(z)^{a_s \left(1 - \frac{\kappa}{\rho}\right)}\right)^{\frac{\rho}{\kappa}}. \quad (65)$$

Using the function $\zeta^{i \rightarrow o}$ in the expression for the wage offer distribution and differentiating the later, we obtain

$$\frac{\partial \Upsilon_s^i(z)}{\partial z} = M^G \left(\frac{v_s^i(z)}{V_s}\right) \left(\Omega^i(z) \frac{\partial \Phi_z(z)}{\partial z} + \Omega^o[\zeta^{i \rightarrow o}(z)] \frac{d\Phi_z[\zeta^{i \rightarrow o}(z)]}{dz}\right). \quad (66)$$

Here as well, ignoring for now the skill complementarity, (56) and (66) together with (57) form a system of differential equations which, subject to the boundary conditions $\underline{w}_s^i(z) = \underline{w}_s$ and $\Upsilon_s^i(z) = 0$, returns the policy functions (w_s^i, v_s^i) and the job ladder ranks Υ_s^θ .

G.4 Algorithm

The algorithm has four levels of iteration. The most inner level solves for the policy functions and wage offer distributions using the system of differential equations obtained in Section G.3 while taking into account the complementarity across skills. The second most inner levels iterate on the aggregate number of vacancy, $\{V_s\}_s$, and the production cutoffs, \underline{z} and \hat{z} . The intermediate levels iterate on the indifference function φ , the reservation wages $\{\underline{w}_s\}_s$, and the productivity cutoff of contractor firms, \underline{z}^C . Finally, the outer level iterates on the market clearing condition to solve for the price of outsourcing p .

⁴³While we cannot guarantee that this assumption is satisfied for all calibrations of the model, we always check numerically whether there exists profitable deviations for \underline{z} and \hat{z} and find that these do not exist.

Differential equations Given an outsourcing price o , the function φ , the reservation wages $\{\underline{w}_s\}_s$, the aggregate vacancies $\{V_s\}_s$ and the productivity cutoffs \underline{z} , \hat{z} and \underline{z}^C , the most-inner loop iterates forward on the differential equations to solve for the policy functions and the wage offer distributions. In particular, we iterate twice on the grid of productivity: one time to solve for the policy functions of good producers that hire their service workers in-house, and a second time to solve for the policy functions of good producers that outsource their service workers.

For the first iteration, we need to first find the initial conditions \underline{w}_1^i and Υ_1^i depending on whether the least productive good producers offer the reservation wage to service workers. To find these boundary conditions, we proceed as follows. For each $z \geq \underline{z}^C$ starting from \underline{z}^C for which we know that $w_1^C(\underline{z}^C) = \underline{w}_1$ and $\Upsilon_1^C(\underline{z}^C) = 0$:

1. Compute the policy functions of the least productive good producer with productivity \underline{z} as if this firm was offering wages $w_1^i(\underline{z}) = w_1^C(z)$ and $w_s^i(\underline{z}) = \underline{w}_s$ for $s > 1$:
 - (a) Compute $n_s^i(\underline{z})$ for each s . For $s > 1$, this is $n_s^i(\underline{z}) = n_s(\underline{w}_s)$. For $s = 1$, this is $n_1^i(\underline{z}) = n_1[w_1^C(z)]$.
 - (b) Compute $\{v_s^i(\underline{z})\}_s$ jointly by solving (57) using a non-linear solver.
 - (c) Compute firm output and marginal products of labor.
2. Check if condition (63) holds at z .
 - (a) If it does, set $w_1^i(\underline{z}) = w_1^C(z)$ and start the following algorithm.
 - (b) If not, continue.
3. Compute $\partial_z w_1^C(z)$ and $\partial_z \Upsilon_1^C(z)$ from (58) and (62) respectively and go back to step 1.

Once the initial conditions $\underline{w}_1^i = w_1^C(\underline{z}^C)$ and $\Upsilon_1^i = \Upsilon_1^C(\underline{z}^C)$ have been found, we can iterate once on the differential equations for the good producers that hire their service workers in-house. The starting point for the wage equivalence functions are $\zeta^{i \rightarrow o}(\underline{z}) = \hat{z}$ and $\zeta^{G \rightarrow C}(\underline{z}) = \underline{z}^C$. Then, for any $z \geq \underline{z}$, given that we know $\{w_s^i(z)\}_s$, $\{\Upsilon_s^i(z)\}_s$ and the wage equivalence functions:

1. Compute $\{n_s^i(z)\}_s$ from the labor supplies.
2. Compute $\{v_s^{Fi}(z)\}_s$ jointly by solving (57) using a non-linear solver.
3. Compute $\zeta^{G \rightarrow C}(z')$ from (60), where $z' = z + \Delta_z$ is the next point on the grid of z , and $R^i(z')$ and $n_1^i(z')$ are obtained from linear extrapolation of $\{n_s^i(z - \Delta_z), n_s^i(z)\}_s$. Similarly, compute $\zeta^{i \rightarrow o}(z')$ from (65).
4. Use (61) to compute $\Upsilon_1^i(z')$. Similarly, for $s > 1$, use (66) to compute $\Upsilon_s^i(z')$
5. Compute wages $\{w_s^i(z')\}_s$ from the wage ODEs (56).

Once this iteration over the z 's is finished, proceed to iterate over the \hat{z} to compute the policy functions for good producers that outsource their service workers. Given that these firms do not hire service workers, the first step to find the lower bounds of $w_1^o(\hat{z})$ and $\Upsilon_1^o(\hat{z})$ is not necessary. The remaining

steps are identical except for step 4: in this iteration, it is not required to compute $\zeta^{G \rightarrow C}$ but it is needed to find the numerical inverse of $\zeta^{i \rightarrow o}$. Once this iteration is finished, we have recovered all the policy functions for the good producers, $\{w_s^\theta(z), v_s^\theta(z)\}_{s \in \{1, \dots, S\}, \theta \in \{i, o\}}$, and the firms' rank on the job ladders, $\{\Upsilon_s^\theta(z)\}_{s \in \{1, \dots, S\}, \theta \in \{i, o\}}$. The fact that we only need two iterations allows for a fast computation of the equilibrium despite its complexity.

Inner iteration Given an outsourcing price o , the function φ , the reservation wages $\{\underline{w}_s\}_s$ and the productivity cutoff of contractor firms \underline{z}^C , the inner iteration solves for the aggregate vacancies $\{V_s\}_s$ and the productivity cutoffs. In particular, $\{V_s\}_s$ need to be consistent with equations (50) and (52), while \underline{z} and $\hat{\underline{z}}$ are given in condition 6 of Definition 1.

Intermediate iteration Given an outsourcing price o , the intermediate iteration solves for the indifference function, φ , the reservation wages, $\{\underline{w}_s\}_s$ and the productivity cutoffs of contractor firms \underline{z}^C . Specifically, given the profit functions π^i and π^o , the function φ is found through numerical inversion of the condition $\pi^i(z) = \pi^o[\varphi(z)]$. Then, from the policy functions $\{w_s^\theta(z), v_s^\theta(z)\}_{s \in \{1, \dots, S\}, \theta \in \{i, o\}}$ and the updated indifference function φ , we compute the wage offer distributions $\{F_s\}_s$ that we then use to update the reservation wages according to (53). Finally, the productivity cutoff of contractor firms is computed as $\underline{z}^C = \underline{w}_1 / (p\tau)$.

Outer iteration The outer iteration solves for the price of outsourcing through the market clearing condition (54). The supply of outsourcing services is computed directly from the contractors' policy functions on $[\underline{z}^C, \hat{\underline{z}}^C]$ and from the good producers' policy functions through $\zeta^{G \rightarrow C}$ for $z \geq \hat{\underline{z}}^C$.

G.5 Estimation of parameters

As described in the main text, the estimation of the parameters is broken down into three steps. The first step inverts some equations of the model to directly estimates $\{\delta_s\}_{s=1}^3$ and $\{\zeta_s\}_{s=1}^3$ from the data (see Section G.6 to derive a closed form expression for the EE rate). The second step sets $\xi = 0.5$. Finally, the third step estimates together the remaining parameters via a minimum distance estimator. The remaining parameters can be divided into two sets. The first set of parameters can be estimated by model inversion given the other parameters. This set consists of the matching efficiency, $\{\mu_s\}_s$, and the unemployment benefits, $\{b_s\}_s$. To recover the matching efficiency, note that, in equilibrium, unemployed workers always accept the job offers that they receive, so that $\lambda_s^U = \text{NE}_s$. Furthermore, in the model, $\lambda_s^U = \mathcal{M}(m_s[u_s + \zeta_s(1 - u_s)], V_s) / m_s[u_s + \zeta_s(1 - u_s)]$. Together, we obtain

$$\mu_s = \text{NE}_s \left(\frac{m_s[u_s + \zeta_s(1 - u_s)]}{V_s} \right)^{1-\xi}.$$

Given V_s , the above equation identifies μ_s . To recover the unemployment benefits, rewrite the expression for the reservation wage (53) as

$$\underline{w}_s = \text{RE}_s \mathbb{E}[w_s] + (k_s^U - k_s) \int_{\underline{w}_s} \frac{1 - F_s(w)}{1 + k_s(1 - F_s(w))} dw,$$

where $\text{RE}_s \equiv b_s/\mathbb{E}[w_s]$ is the replacement rate and is targeted in the estimation. Importantly, the above expression is the only equation in which b_s appears and $\mathbb{E}[w_s]$ is independent of b_s up to \underline{w}_s . Hence, by directly setting RE_s , it is possible to compute the reservation wage without knowing b_s and to recover b_s as a residual,

$$b_s = \underline{w}_s - (k_s^U - k_s) \int_{\underline{w}_s} \frac{1 - F_s(w)}{1 + k_s(1 - F_s(w))} dw.$$

To estimate the second set of parameters, we define the loss function

$$\mathcal{L}(\boldsymbol{\theta}) = \sqrt{\frac{1}{N} \sum_{n=1}^N [h_n(\boldsymbol{\theta}) - \hat{h}_n]^2},$$

where $\boldsymbol{\theta}$ is the vector of parameter to be estimated, $\{\hat{h}_n\}_{n=1}^N$ is the set of empirical moments we are targeting, and $h : \mathbb{R}^N \mapsto \mathbb{R}^N$ maps parameters into simulated moments from our model. The simulated moments are computed as exact analogs of the empirical moments. To compute the simulated moments, we simulate a dataset in the (z, ε) space, projecting onto this space the policy functions found in the (z, \hat{z}) space using linear interpolations.

To find the minimum of \mathcal{L} , we use a gradient descent algorithm. That is, starting from $\boldsymbol{\theta}^0$, we obtain a sequence of parameters $\{\boldsymbol{\theta}^j\}_j$ by iterating on $\boldsymbol{\theta}^{j+1} = \boldsymbol{\theta}^j - \gamma_j \nabla \mathcal{L}(\boldsymbol{\theta}^j)$, where the endogenous step size follows the Barzilai–Borwein method. Namely, for $j > 1$,

$$\gamma_j = \frac{\max\{|\boldsymbol{\theta}^j - \boldsymbol{\theta}^{j+1}|^T \cdot |\nabla \mathcal{L}(\boldsymbol{\theta}^j) - \nabla \mathcal{L}(\boldsymbol{\theta}^{j-1})|, 10^{-3}\}}{\|\nabla \mathcal{L}(\boldsymbol{\theta}^j) - \nabla \mathcal{L}(\boldsymbol{\theta}^{j-1})\|^2}.$$

We impose a maximal step size as in Burdakov et al. (2019) to stabilize the descent. The gradient of the loss function is approximated with central finite difference to maximize accuracy. Given that we use $N = 13$ parameters, the loss function $\mathcal{L}(\boldsymbol{\theta})$ is high-dimensional and we cannot check for the existence of local minima. To avoid those, we first search manually to start the algorithm from a $\boldsymbol{\theta}^0$ with a relatively low loss function, in practice $\mathcal{L}(\boldsymbol{\theta}^0) \in [0.3, 0.5]$. The gradient descent attains its minimum at $\mathcal{L}(\boldsymbol{\theta}^*) = .085$. The gradient descent is implemented in Julia and parallelized over 6 CPUs. The descent is run on a standard laptop and takes about one hour to converge.

G.6 Estimation: expression for the employment-employment transition rate

Omit s indices for simplicity. Our argument requires only that the economy be stationary. Index firms by their wage offer w and thie vacancy decision v . Denote $H(v|w)$ the conditional c.d.f. of vacancies given the wage offer. Then

$$EE = \frac{\lambda^E \iint n(w, v)(1 - F(w))dF(w)H(dv|w)}{\iint n(w, v)dF(w)H(dv|w)}.$$

The integral over $H(dv|w)$ produces the vacancy share of goods producers in the numerator and denominator, and hence drops out. Hence,

$$EE = \frac{\lambda^E \int \frac{(1+k)e}{(1+k(1-F(w)))^2} (1-F(w)) dF(w)}{\int \frac{(1+k)e}{(1+k(1-F(w)))^2} dF(w)} = \frac{\lambda^E \int_0^1 \frac{(1-F)dF}{(1+k(1-F))^2}}{\int_0^1 \frac{dF}{(1+k(1-F))^2}},$$

after changing variables to $F = F(w)$. Both integrals admit closed-form expressions, and thus:

$$EE = \lambda^E \frac{((1+k) \log(1+k) - k)/(k^2(1+k))}{1/(1+k)} = \delta \frac{(1+k) \log(1+k) - k}{k}.$$

G.7 Estimation of counterfactual

The baseline parameters of our model are estimated as if the economy was in 2002. To quantify the effects of the rise of outsourcing between 1996 and 2007, we therefore estimate two outsourcing shocks: a negative outsourcing shock that pushes back the economy to 1996, and a positive outsourcing shock that brings the economy to 2007. To compute both shocks, we estimate the mixture of parameter changes that matches the (rescaled) elasticities presented in Table 2. Specifically, let $o(\boldsymbol{\theta})$ be the outsourcing share and $\mathbf{X}(\boldsymbol{\theta})$ the variables of interest (e.g. log value added, log employment, etc.) under the parameter vector $\boldsymbol{\theta}$. For each shock, the new vector of parameters $\boldsymbol{\theta}'$ as to be such that

$$\begin{aligned} \frac{\mathbf{X}(\boldsymbol{\theta}') - \mathbf{X}(\boldsymbol{\theta})}{o(\boldsymbol{\theta}') - o(\boldsymbol{\theta})} &= \boldsymbol{\gamma}, \\ o(\boldsymbol{\theta}') - o(\boldsymbol{\theta}) &= \Delta_o, \end{aligned}$$

where $\boldsymbol{\gamma}$ are the cross-industry scaled elasticity and Δ_o is the size of the outsourcing shock considered. To implement a fast estimation of the shocks, we solve locally for $\boldsymbol{\theta}'$. Taking a first order approximation around $\boldsymbol{\theta}$, the system above becomes

$$\begin{aligned} [\mathbf{D}\mathbf{X}(\boldsymbol{\theta}) - \boldsymbol{\gamma}\mathbf{D}o(\boldsymbol{\theta})] (\boldsymbol{\theta}' - \boldsymbol{\theta}) &= \mathbf{0}, \\ \mathbf{D}o(\boldsymbol{\theta})(\boldsymbol{\theta}' - \boldsymbol{\theta}) &= \Delta_o, \end{aligned} \tag{67}$$

where \mathbf{D} is the differential operator with respect to $\boldsymbol{\theta}$ and $\mathbf{D}\mathbf{X}$ and $\mathbf{D}o$ can be computed by perturbing the economy around its 2002 steady state. We can then invert this linear system to recover $\{\boldsymbol{\theta}_{1996}, \boldsymbol{\theta}_{2007}\}$. To construct the counterfactual reported in Section 5, we fit cubic polynomials on $\{\boldsymbol{\theta}_{1996}, \boldsymbol{\theta}_{2002}, \boldsymbol{\theta}_{2007}\}$ for each dimension of $\boldsymbol{\theta}$. We then use these polynomials to infer a vector of shock for each year in the 1996-2007 time frame. We also use the fitted cubic polynomials to extrapolate the consequences of outsourcing till 2016.

G.8 Accounting

This section details how the main micro and macro variables are computed in the model. For that, suppose that we have simulated a cross-sectional data set at the firm level from our model. Let i and

j describes the identity of a firm in this data set.

Goods producers profits and value added. Profits of goods producer i are

$$\text{Profits}_i^G = R_i^G - \sum_s w_{is}^G n_{is}^G - p\varepsilon_i n_i^G - \sum_s c(v_{is}^G) - \eta,$$

where the notation follows closely that of Section G.1. The value added of goods producer i is revenues net of spending on intermediaries, or

$$\text{VA}_i^G = R_i^G - p\varepsilon_i n_i^G.$$

The variable ε_i indeed represents iceberg costs faced by a given goods producer i . Goods producer thus i needs to purchase $\varepsilon_i n_i^G$ units of labor in the labor service market to obtain n_i^G units of effective labor in production.

Contractor profits and value added. Contractors j make profits

$$\text{Profits}_j^C = R_j^C - w_j^C n_j^C - c(v_j^C) = p\tau z_j n_j^C - w_j^C n_j^C - c(v_j^C),$$

and have value added

$$\text{VA}_j^C = R_j^C = p\tau z_j n_j^C.$$

Aggregate output. Aggregate output is the sum of value added of all sectors of the economy. Aggregate output coincides with the amount of goods available for consumption for workers, who receive wage payments, and capital owners, who receive vacancy costs and fixed costs. Thus,

$$\begin{aligned} \text{Ag. output} &= \sum_i \text{VA}_i^G + \sum_j \text{VA}_j^C \\ &= \sum_i \left(R_i^G - p\varepsilon_i n_i^G \right) + p \sum_j \tau z_j n_j^C \\ &= \sum_i \left(R_i^G - p\varepsilon_i n_i^G \right) + p \sum_i \varepsilon_i n_i^G \\ &= \sum_i R_i^G \end{aligned} \tag{68}$$

where the first equality uses the definitions of value added, and the second equality uses labor services market clearing (54).

Aggregate TFP We define TFP as

$$\text{TFP} = \frac{\text{Ag. output}}{\bar{N}}, \tag{69}$$

where the labor aggregator \bar{N} is defined as

$$\bar{N} = \left(\prod_s \bar{N}_s^{a_s} \right)^\rho,$$

and $\bar{N}_s = m_s(1 - u_s)$ is employment of skill s . To capture reallocation towards less productive contractors, we define the adjusted aggregator

$$\tilde{N} = \tilde{N}_1^{\rho a_1} \left(\prod_{s=2}^3 \bar{N}_s^{a_s} \right)^\rho, \quad \tilde{N}_1 = \bar{N}_1^G + \tau_1 \bar{N}_1^C,$$

where \bar{N}_1^G denotes aggregate employment of skill 1 by goods producers, and \bar{N}_1^C aggregate employment by contractors. The effective measure \tilde{N}_1^G encodes the effective amount of labor used for task 1 in the economy. The ratio

$$\frac{\tilde{N}}{\bar{N}} = \left(\frac{\tilde{N}_1}{\bar{N}_1} \right)^{\rho a_1} = (\tau_1 x_1^C + (1 - x_1^C))^{\rho a_1},$$

where $x_1^C = \bar{N}_1^C / \bar{N}_1$ is the employment share of contractors among low skill service workers, captures the TFP effect of reallocation towards more or less productive contractors. When $\tau_1 < 1$ and x_1^C rises as outsourcing increases, \tilde{N}/\bar{N} decreases: workers are reallocated towards less productive jobs as far as production of labor services is concerned. Aggregate TFP then writes

$$\text{TFP} = \frac{\text{Ag. output}}{\tilde{N}} \times \frac{\tilde{N}}{\bar{N}},$$

and so changes in aggregate TFP are

$$\Delta \log \text{TFP} = \underbrace{\Delta \log \frac{\text{Ag. output}}{\tilde{N}}}_{\substack{\text{Allocative efficiency} \\ \text{given effective labor} \\ \text{in the economy}}} + \underbrace{\Delta \log \frac{\tilde{N}}{\bar{N}}}_{\substack{\text{Productivity gains/losses} \\ \text{from contractor comparative} \\ \text{advantage/disadvantage:} \\ \text{change in effective labor}}}.$$