Outsourcing, Inequality and Aggregate Output*

Adrien Bilal — *Stanford University*

Hugo Lhuillier — *University of Chicago*

February 27, 2025

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. More productive firms select into outsourcing to save on labor costs and higher wage premia. Outsourcing leads firms to raise output and labor demand. Contractor firms pay lower wages. We find reduced-form support for all three implications in French administrative data, instrumenting revenue productivity with export demand shocks and outsourcing costs using variation in occupational exposure. After proving identification and structurally estimating the model, we find that the emergence of outsourcing in France lowers low skill service worker earnings and welfare by 3.1% but raises aggregate output by 1.8%.

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

*Adrien Bilal: Stanford University, CEPR and NBER. adrienbilal@stanford.edu. Hugo Lhuillier: University of Chicago. hugolhuillier@uchicago.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 (Dube and Kaplan, 2010; Goldschmidt and Schmieder, 2017). This outsourcing wage penalty suggests that rising domestic outsourcing redistributes away from workers. At the same time, firms may scale up more efficiently by contracting out, leading to aggregate productivity and employment gains that benefit workers. Despite the prevalence of outsourcing in the labor market, the tension between its distributional and aggregate 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. This view suggests that outsourcing raises aggregate output through Total Factor Productivity (TFP) as workers are reallocated toward more efficient contractors. Another perspective, the *cost-saving view*, holds that contractor firms enable their clients to sidestep costly hiring and save on labor expenditures. This view has more nuanced implications, as the reallocation of workers toward less efficient contractors may be 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 determinants 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 consists of labor market frictions and 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, and wages become an effective hiring tool. As a result, firms face an upward-sloping labor supply curve by skill: higher pay lets them attract and retain more workers from competitors. We characterize wage and employment distributions in closed form. More productive firms with a larger target size pay higher wages, and wage inequality emerges in equilibrium.[1]

We then introduce contractor firms in our environment. Contractor firms hire workers in the same frictional labor markets and with the same recruiting technology as goods producers. Instead of producing a consumption good, contractor firms sell the labor services of their employees at an

[1] We use "wages" for concreteness but our theory applies equally to total compensation inclusive of benefits.

equilibrium price in a competitive labor service market. In the aggregate, contractor firms effectively expand managerial resources available for recruiting.

Consistent with the cost-saving view, goods producers may now 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. Consistent with the comparative advantage view, contractor firms may be more or less productive at generating labor services than goods producers, for instance because of gains from specialization or communication and coordination costs from 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 attain 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 rise while the marginal product of labor falls. All else equal, declines in outsourcing costs thus increase 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 strong 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 micro-foundation of firm-specific labor supply curves implies that this ranking of wages translates into a ranking of worker flows across firms.

Our theory thus helps discriminate between the comparative advantage and cost-saving views. The comparative advantage view is difficult to reconcile 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.

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 guarantee existence and uniqueness. By contrast, our analysis requires productivity heterogeneity and decreasing returns to maintain a well-defined distribution of firm size when goods producers outsource. 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 and a lower bound on firm productivity that precludes non-smooth equilibria.

In the second part of the paper, we test the implications of our theory using administrative data from France. 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 and service workers using occupation 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, an increase that is mirrored in employment at contractor firms. Extrapolating beyond 2007 suggests that outsourcing keeps rising through 2016.

We first examine the distributional effect of outsourcing. Consistent 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 in French data. Contractor firms pay wages that are 14% below other firms after controlling for individual worker heterogeneity. In line with our microfoundation of the labor supply curve, we show that contractors also hire less from employment than other firms, churn more through workers, and have negative net poaching.

Second, we document that firms select into outsourcing. We find a robust correlation between the outsourcing share—outsourcing expenditures relative to total expenditures on labor—and revenues or value added across and within firms. Since this relationship may be endogenous to unobserved changes in outsourcing costs such as improvements in information technologies, we isolate the effect of a change in revenue productivity by instrumenting revenues 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; Borusyak et al., 2021). Our Two Stage Least Squares (2SLS) estimate implies that an increase in revenue productivity that implies a 10% increase in revenues leads to a 0.33 percentage points (p.p.) rise in the outsourcing share.

Third, we provide evidence for the productivity effect. We isolate declines in firm-level outsourcing costs by instrumenting the outsourcing share with shift-share outsourcing cost shocks using variation in exposure at the firm level (Goldsmith-Pinkham et al., 2020): we interact firm-level initial occupation shares with changes in average outsourcing expenditures at the occupation level. Our 2SLS estimate indicates that a decline in outsourcing costs that imply a 1 p.p. rise in the outsourcing share leads to a 9% rise in revenues.

We evaluate whether alternative mechanisms could explain our findings. We do not find strong support that demand volatility, equity concerns or union wage-setting or size-dependent policies are more important drivers of outsourcing than firm scale.

In the third part of the paper, we develop and structurally estimate an extended framework suitable for quantification. Firms now face flexible curvature in job creation cost functions that allow them to hire in-house without raising wages too rapidly. Contractors have a possible comparative advantage in job creation to match their relative size. Firms enter freely subject to entry costs.

We provide a constructive proof of identification for 21 of the 25 parameters of the model by extending the strategy in Bontemps et al. (2000) to our setting. We estimate the remaining 4 with a Method of Simulated Moments (MSM) estimator. 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. While our framework can accommodate high skill outsourcing, we purposefully focus on low skill outsourcing for whom the equity-efficiency trade-off is particularly salient.

The estimated model highlights that contractors specialize in recruiting activities rather than in production activities. Observed outsourcing expenditure and employment shares imply that workers at contractor firms are 50% less efficient than in-house workers. This result falsifies the comparative advantage view but is consistent with communication and monitoring frictions. Given relative size and wages across firms, contractors are more efficient at recruiting workers than goods producers: their core activity is to recruit workers.

We then quantify the race between the distributional and aggregate effects of outsourcing with two counterfactuals. The first is an outsourcing ban. The second investigates how the rise in outsourcing between 1997 and 2016 due to a decline in entry costs reshaped the labor market in France.

We find that low skill service workers lose from outsourcing: they are better off under the outsourcing ban, and worse off because of the rise in outsourcing between 1997 and 2016. Expected earnings—which coincide with welfare in our environment—of low skill service workers rise by 3.1% under the outsourcing ban, and fall by 1.7% between 1997 and 2007.

In partial equilibrium, outsourcing reallocates service workers toward low-paying contractor firms: their baseline employment share is 19% in 1997, and rises by 21 p.p. by 2007. Together with the outsourcing wage penalty, this reallocation implies that service workers gain 2.5% in expected wages under an outsourcing ban relative to 1997. The subsequent reallocation between 1997 to 2007 reduces earnings by 2%.

General equilibrium channels partly offset this decline. The productivity effect leads to substantial employment gains for low skill service workers. Under an outsourcing ban, these employment gains do not materialize, and earnings are 2% lower. Between 1997 and 2007, rising employment leads to a 2.5% increase in earnings, and 5% by 2016. An additional competition channel further reverses some of the earnings losses between 2007 and 2016. As outsourcing keeps rising, wages at low-paying firms start increasing as competition intensifies at the bottom of the job ladder. Our decomposition highlights that reduced-form approaches that only pick up the first, partial equilibrium impact miss key general equilibrium adjustments and are a partial representation of welfare losses of service workers.

In the aggregate, the economy benefits from outsourcing due to the productivity effect. An outsourcing ban lowers output by 1.8% relative to 1997. This reduction is due to a combination of employment and Total Factor Productivity (TFP) contributions. Employment is higher when outsourcing is allowed because it lets the economy tap into additional managerial resources for job creation. Outsourcing also increases allocative efficiency by increasing employment at high marginal product of labor firms that were constrained by labor market frictions, raising TFP. However, outsourcing simultaneously reallocates workers toward less efficient contractors, which puts downward pressure on TFP. We find this offsetting channel to be quantitatively important. By contrast, the comparative advantage view would imply that both effects are always positive, crucially overstating aggregate TFP gains.

Aggregate output gains from outsourcing flow back to core low skill and high skill workers, who lose 1.7% in earnings from an outsourcing ban and gain nearly 2% between 1997 and 2007. The estimated revenue function implies that most of these gains accrue through higher wages and stable markdowns rather than employment gains or changes in monopsony power. Thus, between-skill inequality rises. Overall, our analysis thus indicates that outsourcing is detrimental to low skill service workers while it benefits other workers and leads to aggregate output gains.

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. (2023) document that domestic outsourcing is increasing and that outsourced workers experience wage declines in the U.S., Germany and Argentina.[2] Abraham and Taylor (1996) provide an early discussion of the cost-saving and comparative advantage views. Consistent with the productivity effect, Bertrand et al. (2020) show that an increase in the supply of contract labor helped Indian firms scale up. Felix and Wong (2024) find market-level employment gains following outsourcing in Brazil. We contribute to this literature by providing the first unified general equilibrium theory of outsourcing in which we disentangle the comparative advantage and cost-saving views, tie the productivity and distributional effects together, and analyze the trade-off between both forces in the aggregate.[3]

Second, 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 and Mortensen, 1998; Berg and Ridder, 1998; Bontemps et al., 2000; Card, Cardoso, et al., 2018; Sorkin, 2018; Engbom and Moser, 2022; Heise and Porzio, 2022; Gouin-Bonenfant, 2022; Berger et al., 2022; Lhuillier, 2022; Lamadon et al., 2022; Morchio and Moser, 2024) by departing from constant returns to scale and perfect substitutability between workers when the labor supply curve is an equilibrium object and there is productivity dispersion. We provide sufficient conditions to solve for an equilibrium in this setting, which has remained an open problem since at least Mortensen and Vishwanath (1991).

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.[4]

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 in the Online Appendix.

[2]See also Segal and Sullivan (1997), Katz and Krueger (2017), and Katz and Krueger (2019) regarding alternative work arrangements in the U.S. and Muñoz (2023) for the related role of posted workers in the European Union.

[3]Giannoni and Mertens (2019) emphasize the impact of outsourcing on the labor share in the U.S. Bergeaud et al. (2025) 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. Since this paper was first circulated, Spitze (2022) developed a related model with bargaining in which outsourcing reduces wages at the aggregate level but does not match the empirical outsourcing wage penalty between firms without further assumptions. Micco and Perez (2024) develop a framework with exogenous differences in bargaining power—and hence wages—between permanent and temporary workers. Bostanci (2022) studies the trade-off between trade secrets and the productivity effect of outsourcing.

[4]A related literature studies the make-or-buy choice of firms (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. This restriction allows us to focus on the labor market and aggregate consequences of domestic outsourcing.

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.[5] 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 $[\underline{z}, \overline{z}]$, $\underline{z}, \overline{z} \in (0, +\infty)$. The corresponding cumulative distribution function Γ admits a finite and continuous density. We assume that $\underline{z}$ is large enough relative to $\sup_s b_s$ to ensure that the equilibrium is well-behaved. A firm with productivity z that hires a measure n_s of workers of each skill s generates revenue $R(z, \boldsymbol{n})$, where $\boldsymbol{n} = \{n_s\}_s$ denotes the vector of employment across worker types. R is twice continuously differentiable and increasing in each argument.

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 intensity 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 Labor supply, wages and employment

The skill-specific labor supply curve faced by each firm depends on the job search behavior of workers. As in Burdett and Mortensen (1998), workers maximize their income in equilibrium and switch jobs when they receive wage offers above their current pay drawn from an equilibrium distribution $F_s(w)$. The movement of workers along the job ladder implies that the number $N_s(w)$ of employed workers per wage offer w for every skill s is:

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

where $e_s = \frac{\lambda_s^U m_s}{\delta_s + \lambda_s^U}$ is the measure of employed workers of skill s, $k_s = \frac{\lambda_s^E}{\delta_s}$, and $F_s(w^-)$ denotes the left-limit of F_s at w.[6] The derivations mirror Burdett and Mortensen (1998) and are detailed in Online Appendix D.1 for completeness.

The labor supply curve $n_s(w)$ is the number of workers per firm posting wage w. It is related to the

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

[6] For any function f and real argument x, we denote by $f(x^+) \equiv \lim_{y>0, y\downarrow 0} f(x+y)$ and $f(x^-) \equiv \lim_{y>0, y\downarrow 0} f(x-y)$ the right- and left-limit of a function f around its argument x.

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. This labor supply curve $n_s(w)$ is non-decreasing in the wage w. Firms who offer higher wages can retain a larger fraction of their workforce while poaching more workers from their competitors, thus attaining a larger size. The labor supply elasticity depends on the skill-specific distribution of wage offers in the economy, $F_s(w)$, which is determined in equilibrium.

Firms choose wage offers $\{w_s(z)\}_s$ to maximize flow profits when the discount rate converges to zero as in Burdett and Mortensen (1998). We also operate under this limit.[7] Flow profits are given by:

$$\pi(z) \quad = \quad \max_{\{w_s, n_s\}_s} R\big(z, \{n_s\}_s\big) - \int w_s n_s ds \quad \text{subject to} \quad n_s \leq n_s(w_s), \tag{2}$$

where firms take as given their skill-specific labor supply curves.

Unless the wage offer distribution $F_s(w)$ can be characterized more precisely, the problem in equation (2) is intractable in general equilibrium. The distribution $F_s(w)$ is the solution of a complex fixed point problem: wage-setting strategies determine equilibrium wages but also depend on the wage offer distribution itself through labor supply curves.

The wage-posting literature—from Burdett and Mortensen (1998) to Engbom and Moser (2022)—has recognized this challenge and instead leveraged a key simplifying assumption to make progress. Under constant returns and perfect substitutability of workers in production, $R(z, \boldsymbol{n}) = z \int n_s ds$, the problem (2) can be split at the match level. Once decoupled across matches and under constant returns, (2) exhibits positive cross-derivatives in productivity and employment regardless of size. As a result, wages are increasing in productivity z, and the wage offer distribution at equilibrium wages coincides with the exogenous productivity distribution, $F_s(w_s(z)) = \Gamma(z)$.

However, to introduce outsourcing, we require a well-defined boundary of the firm through decreasing returns to scale in production. Otherwise, firm size is not well-defined once firms have the option of hiring labor services in a competitive market. Quantitatively, possible interactions between workers in production are also key to understand the distributional implications of outsourcing. Handling these features together with productivity heterogeneity has remained an open problem since Mortensen and Vishwanath (1991) who pointed out that the usual arguments for monotone wages and uniqueness may 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). $(z, \boldsymbol{n}) \mapsto R(z, \boldsymbol{n})$ 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 propor-

[7] We derive the formulation in equation (2) from the dynamic problem of the firm in Online Appendix D.8.

tions. 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, Online Appendix D.3 shows that Assumption (A) is compatible with a CES production function between worker types coupled with CES demand between firm varieties leading to decreasing returns in revenues, as long as varieties are more substitutable than workers.

We impose Assumption (A) in the remainder of this paper. We first show that Assumption (A) guarantees that any equilibrium with a smooth offer distribution has a simple structure. Proposition 1 is our first main theoretical result. It shows that supermodularity is sufficient to generalize the logic behind the sharp equilibrium characterization in Burdett and Mortensen (1998).

Proposition 1. *(Wages and employment)*
In any equilibrium with a continuous wage offer distributions $F_s(w)$, wages $w_s(z)$ are strictly increasing and continuous in firm productivity z, and satisfy $F_s(w_s(z)) = \Gamma(z)$. Wages $w_s(z)$ and the number of workers of skill s hired by firm z, $n_s(z) = n_s(w_s(z))$, satisfy:

$$n_s(z) = \frac{(1+k_s)e_s}{M^g[1+k_s(1-\Gamma(z))]^2} \qquad w_s(z) = \underline{w}_s \frac{n_s(\underline{z})}{n_s(z)} + \int_{\underline{z}}^{z} \frac{\partial R}{\partial n_s}(x, \boldsymbol{n}(x)) \frac{n_s'(x)dx}{n_s(z)}.$$

Proof. See Appendix A.1. □

Firm size in Proposition 1 depends only on the ranking of firms, $\Gamma(z)$, because firm size is fully determined by worker flows up the job ladder.[8] Competition for workers along the job ladder underpins the wage equation in Proposition 1. Productive firms raise their wages to poach workers from lower-productivity firms to reach their target size. The 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 Online Appendix D.2, and integrating up to productivity z. The resulting wage function is a weighted average of the marginal product of labor at lower-ranked firms. Proposition 1 nests the wage equations with linear revenue in e.g. Burdett and Mortensen (1998) and Engbom and Moser (2022).

Proposition 1 requires an equilibrium with a continuous wage offer distribution to exist. As discussed by Mortensen and Vishwanath (1991) however, wage-posting models with decreasing returns to scale can exhibit either no smooth equilibria or multiple coordination equilibria due to the emergence of mass points.[9] We overcome these difficulties with two observations.

The first observation ensures existence and uniqueness among smooth equilibria by leveraging our assumption that $\underline{z}$ is large enough relative to $\sup_s b_s$. This assumption together with mild technical regularity conditions ensures that decreasing returns do not equalize the marginal product of labor across firms, in which case mass point may emerge. Existence and uniqueness then follows from a standard guess and verify argument leveraging Proposition 1. We formalize our technical regularity conditions in Online Appendix D.4 and call them Assumption (A').

[8] We 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.

[9] If 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.

The second observation is that coordination equilibria are unstable when a smooth equilibrium exists. We consider a trembling-hand equilibrium refinement concept that overcomes possible multiplicity. If firms make small mistakes in wage-setting, no mass point can arise. When dispersion in mistakes vanishes asymptotically so that we recover the maximization problem in (2), the only equilibrium that survives has a smooth wage distribution, is unique and is described in Proposition 1. We formalize our trembling-hand refinement in Online Appendix D.4 and call it Assumption (B).

Proposition 2. *(Existence and uniqueness)*
Under Assumption (A'), there exists an equilibrium with a continuous wage offer distribution and is described in Proposition 1. This equilibrium is unique among equilibria with a continuous wage offer distribution. Under Assumption (B), this equilibrium is unique among all possible equilibria.

Proof. See Online Appendix D.4. □

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. Ownership of contractors is separate from ownership of goods producers.

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.[10] Consistent with the comparative advantage view, contractors may be better or worse than goods producers at producing labor services. We capture this possible comparative advantage through a productivity wedge $\tau_s \lesseqgtr 1$ relative to goods producers. Contractor firms thus solve the profit-maximization problem:

$$\pi_s^C = \max_w \big(\tau_s p_s - w\big) n_s(w). \tag{3}$$

Despite being homogeneous, contractors offer heterogeneous wages between which they are indifferent in equilibrium as in Burdett and Mortensen (1998).

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

We propose three simple micro-foundations for our measure of comparative advantage τ_s, detailed in Online Appendix D.5. 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 the inverse of an iceberg trade cost between contractor firms and goods producers that captures the idea that communication, monitoring and coordination may be more difficult between than within firms. In our second micro-foundation, contractor firms combine a small amount of capital and labor according to a Cobb-Douglas production function, and 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, selection and productivity effect

Consistent with the cost-saving view, goods producers may now rent labor services from contractors or hire workers in-house in a frictional labor market. The hire-or-rent decision of good producers is subject to idiosyncratic outsourcing costs: to rent n_s labor services, goods producers must spend $p_s\varepsilon_s n_s$, 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. While these outsourcing costs do not materially affect the comparative statics in this section, they highlight the key endogeneity challenges we face when we confront our theory with the data in Sections 3.4 and 3.5.[11]

Goods producers solve:

$$\pi(z,\boldsymbol{\varepsilon}) = \max_{\substack{\{n_s\}_s,\{w_s\}_s,\\ \{o_s\}_s\in\{0,1\}^S}} R\big(z,\{n_s\}_s\big) - \int \Big[(1-o_s)w_s + o_s p_s \varepsilon_s\Big] n_s ds \quad \text{s.t.} \quad n_s \leq n_s(w_s) \text{ if } o_s = 0. \tag{4}$$

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 $o_s = 1$. We denote the vector of outsourcing costs by $\boldsymbol{\varepsilon} \equiv \{\varepsilon_s\}_s$.[12] In what follows we always consider an equilibrium in which there is some outsourcing.

If the goods producer hires in-house ($o_s = 0$), it faces an upward-sloping labor supply curve $n_s(w)$. 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 to save on labor costs.

However, when a goods producer is unproductive and targets a small size n_s, it moves down its

[11] For the problem of goods producers to be well-defined under outsourcing, we further require that R is strictly concave in $\{n_s\}_s$ and satisfies a strict Inada condition: that there exists a constant $\rho \in (0,1)$ such that for all $s, z, \boldsymbol{n}$, $0 < \frac{n_s \partial R(z,\boldsymbol{n})/\partial n_s}{R(z,\boldsymbol{n})} \leq \rho < 1$. Importantly, this assumption is weak. It simply ensures that the elasticity of the revenue function is bounded away from below one, so that demand for outsourced labor is well-defined. This strict Inada condition is satisfied for a wide class of revenue functions, for instance in the nested CES example that we describe in Online Appendix D.3 or for the Cobb-Douglas revenue function that we use in Sections 4 and 5.

[12] For simplicity of exposition, we require that goods producers do not mix in-house and outsourced employment for a given skill s. If goods producers could mix, they would. Our results are not materially affected if we lift this restriction.

supply curve. The price of outsourcing p_s then exceeds in-house wages since it reflects both the wage paid to employees of contractor firms as well as compensation to contractors. Outsourcing is then less attractive than hiring in-house.

Thus, as productivity rises conditional on outsourcing costs, goods producers demand more outsourced workers and outsourcing expenditures rise. Proposition 3 formalize this discussion and characterizes selection into outsourcing.

Proposition 3. *(Selection into outsourcing)*
There exists a threshold productivity function $\hat{z}_s(\boldsymbol{\varepsilon})$, such that outsourcing of skill s occurs if and only if $z \geq \hat{z}_s(\boldsymbol{\varepsilon})$. The threshold $\hat{z}_s(\boldsymbol{\varepsilon})$ is increasing in 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.2. □

In addition, outsourcing lets goods producers expand their activities because it lowers labor costs, incentivizing them to use more labor services. A reduction in outsourcing costs conditional on productivity for the marginal firm leads goods producers to increase revenues and labor demand up to their unconstrained scale, and lower their marginal product of labor. Proposition 4 formalizes this effective productivity effect of outsourcing.

Proposition 4. *(Productivity effect)*
Labor demand is larger for firms that outsource just to the right of the outsourcing threshold, relative to firms that operate in-house just to its left: $n_s(\hat{z}_s(\boldsymbol{\varepsilon})^+, \boldsymbol{\varepsilon}) > n_s(\hat{z}_s(\boldsymbol{\varepsilon})^-, \boldsymbol{\varepsilon})$. As a result, revenues are also larger and the marginal product of labor is smaller: $R(\hat{z}_s(\boldsymbol{\varepsilon})^+, \boldsymbol{n}(\hat{z}_s(\boldsymbol{\varepsilon})^+, \boldsymbol{\varepsilon})) > R(\hat{z}_s(\boldsymbol{\varepsilon})^-, \boldsymbol{n}(\hat{z}_s(\boldsymbol{\varepsilon})^-, \boldsymbol{\varepsilon}))$ and $R_{n_s}(\hat{z}_s(\boldsymbol{\varepsilon})^+, \boldsymbol{n}(\hat{z}_s(\boldsymbol{\varepsilon})^+, \boldsymbol{\varepsilon})) < R_{n_s}(\hat{z}_s(\boldsymbol{\varepsilon})^-, \boldsymbol{n}(\hat{z}_s(\boldsymbol{\varepsilon})^-, \boldsymbol{\varepsilon}))$.

Proof. See Appendix A.3. □

Together, Propositions 3 and 4 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 and labor market frictions, relative to a competitive labor market. Outsourcing helps firms bypass those managerial constraints and improves the allocation of labor in the economy.

2.3 The distributional effect

At the same time as it reallocates workers toward goods producers with a high marginal product of labor, outsourcing changes the equilibrium wage distribution. How wages react to outsourcing in turn depends on whether contractors have a comparative advantage or disadvantage.

Proposition 5. *(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 $w_s^{\text{cont.}}$, $w_s^{\text{cont.}} < w_s(\hat{z}_s(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon}) \leq p_s \varepsilon_s$.*

(b) *Suppose that only workers of skill $s = 1$ can be outsourced. There exists a threshold $\overline{\tau}_1 > 1$ such that for all $\tau_1 \geq \overline{\tau}_1$, contractors pay higher wages than the marginal goods producer: for any contractor wage $w_1^{\text{cont.}}$, $w_1^{\text{cont.}} > w_1(\hat{z}_1(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon})$, with $w_1(\hat{z}_1(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon}) \leq p_1 \varepsilon_1$.*

Proof. See Appendix A.4. □

To understand the forces at work in Proposition 5, consider the neutral case $\tau_s = 1$ first. In that case, the outsourcing industry can only be profitable in equilibrium if contractors offer lower wages than in-house goods producers that are close to indifference with outsourcing. Specifically, 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 contractors firms is below the marginal product of labor at in-house goods producers at $\hat{z}_s$. Ranking marginal products of labor lets us rank wages.

Technological disadvantage $\tau_s < 1$ only widens the gap between the marginal product of labor of contractors relative to the marginal outsourcing firm at $\hat{z}_s$. When contractors instead have a large enough comparative advantage $\tau_1 \geq \overline{\tau}_1$, their marginal product outstrips that of goods producers and so they pay higher wages. Regardless of the particular value of comparative advantage or disadvantage, outsourcing always removes the best-paying opportunities in the labor market. Because hiring at wages above the price of outsourcing is never profitable, in-house wages are capped by $p_s \varepsilon_s$.

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 toward high-paying firms. Thus, the wage ordering between contractors and goods producers from Proposition 5 translates immediately into an ordering of several other labor market statistics. We define the fraction of hires from employment of any firm paying wage w as:

$$\mathrm{HE}_s(w) = \frac{q_s(1-\phi_s)G_s(w)}{q_s\big[\phi_s + (1-\phi_s)G_s(w)\big]}, \tag{5}$$

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:

$$\mathrm{Churn}_s(w) = \frac{\big[\delta_s + \lambda_s^E(1 - F_s(w))\big] n_s(w)}{n_s(w)}. \tag{6}$$

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:

$$\mathrm{NP}_s(w) = \frac{q_s(1-\phi_s)G_s(w) - \lambda_s^E(1-F_s(w))n_s(w)}{n_s(w)}. \tag{7}$$

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). Corollary 1 characterizes how contractors compare to goods producers along all those labor market statistics.

Corollary 1. *(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: $\mathrm{HE}_s^{\mathrm{cont.}} < \mathrm{HE}_s(\hat{z}_s(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon})$*,* $\mathrm{Churn}_s^{\mathrm{cont.}} > \mathrm{Churn}_s(\hat{z}_s(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon})$ *and* $\mathrm{NP}_s^{\mathrm{cont.}} < \mathrm{NP}_s(\hat{z}_s(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon})$*.*

(b) Suppose that only workers of skill $s = 1$ can be outsourced. There exists a threshold $\overline{\tau}_1 > 1$ such that for all $\tau_1 \geq \overline{\tau}_1$, contractors hire more from employment, have lower churn and higher net poaching than the marginal goods producer: $\mathrm{HE}_1^{\mathrm{cont.}} > \mathrm{HE}_1(\hat{z}_1(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon})$*,* $\mathrm{Churn}_1^{\mathrm{cont.}} < \mathrm{Churn}_1(\hat{z}_1(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon})$ *and* $\mathrm{NP}_1^{\mathrm{cont.}} > \mathrm{NP}_1(\hat{z}_1(\boldsymbol{\varepsilon}), \boldsymbol{\varepsilon})$*.*

Proof. See Appendix A.5. □

Proposition 5 and Corollary 1 capture the distributional effect of outsourcing. When a goods producer at $\hat{z}_s$ decides to outsource because its idiosyncratic outsourcing cost drops 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 partial equilibrium, the labor reallocation that follows a rise in outsourcing depresses earnings by moving workers away from the highest-paying goods producers and toward low-paying contractor firms.

Proposition 5 highlights that the outsourcing wage penalty helps disentangle the comparative advantage view from the cost-saving view. Consistent with the comparative advantage view, our model nests the possibility that firms outsource because contractor firms have a strong comparative advantage in producing services when $\tau_1 > \overline{\tau}_1$, rather than to save on costs per worker. In this case, Proposition 5 suggests 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. (2023) 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 5 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.

Of course, Proposition 5 leaves open an intermediate region $\tau_1 \in [1, \overline{\tau}_1]$ for which contractors have a comparative advantage in producing labor services, but still pay lower wages than marginal goods producers who outsource. This ambiguous intermediate region requires us to estimate a quantitative version of our model in Section 4 to fully disentangle the comparative advantage view from the cost-saving view.

Together, Propositions 3 to 5 and Corollary 1 characterize the key tension between the productivity and distributional effects of outsourcing. Before we test them empirically, the next section describes how these forces shape the equilibrium wage distribution.

Figure 1: Labor supply and wage distributions in equilibrium.

(a) Labor supply.

(b) Wage distributions.

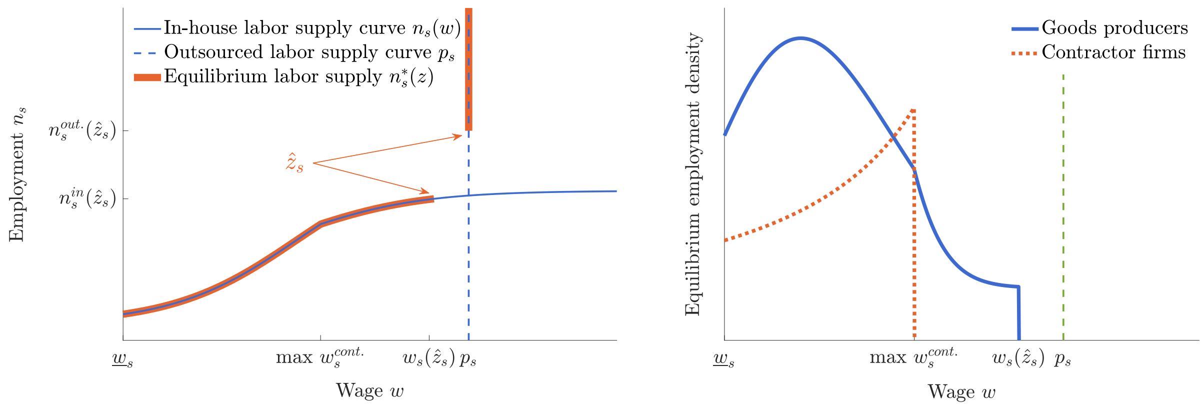

Note: Panel (a): in-house, outsourced and equilibrium labor supply curves of goods producers. Panel (b): equilibrium wage distributions of contractors and goods producers.

2.4 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:

$$M^g \int \int_{\hat{z}_s(\boldsymbol{\varepsilon})}^{\overline{z}} \varepsilon_s n_s(z, \boldsymbol{\varepsilon}) d\Gamma(z, \boldsymbol{\varepsilon}) = \tau_s M_s^c \int n_s^{\text{cont.}}(w) dF_s(w). \tag{8}$$

With some outsourcing, the equilibrium has several regions. For brevity, we focus on the case without idiosyncratic outsourcing costs $\varepsilon_s \equiv 1$. We describe the simplest structure of the equilibrium in this section, and provide a full description in Online Appendix D.6. 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 $\hat{z}_s$ and wage $w_s(\hat{z}_s)$. Proposition 4 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 5 does not restrict the size of the gap between the maximal in-house wage and the outsourcing price beyond $w_s(\hat{z}_s) \leq p_s$. 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 $w_s(\hat{z}_s)$.

Equipped with our characterization of equilibrium, Propositions 3 to 5 and Corollary 1, we now confront our theory with the data.

3 Reduced-form evidence

This section starts by describing our data. Next, we discuss aggregate trends in outsourcing in France. We then test our three main predictions: the distributional effect, selection into outsourcing, and the productivity effect. Finally, we discuss 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 observe employment, payroll, sales and purchases of intermediate inputs. We construct value added as sales minus purchases of intermediate inputs. 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. The EAE is stratified by sector. For instance in services, the EAE surveys large firms with more than 30 employees or sales above 5 million euros. Small firms are surveyed every two years.

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 (*Données 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 can link these four data sources together using a common firm tax identifier. For most of our firm-level exercises, we focus on a firm-level dataset in which we link FICUS, EAE, DADS *Postes* and the Trade data. In this firm-level dataset, we aggregate years into three periods 1997-1999, 2000-2002, 2003-2007 and keep only firms with at least one in-house employee, positive wages and value added to limit measurement error due to non-response and alternating sampling in the EAE. 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.[13] Our final firm-level dataset consists of 216,051 firm-periods, although the underlying data has more than 500,000 firm-year pairs. For most of our worker-level exercises, we focus on a worker-level dataset that only includes the DADS *Panel*.

Table 7 in Online Appendix E presents summary statistics of our firm-level dataset (FICUS-EAE-DADS *Postes*-Trade). Due to the sampling in the EAE survey, our sample covers firms that are larger than the typical French firm on average. Table 8 in Online Appendix E presents summary statistics of our worker-level dataset (DADS *Panel*), which is representative of the French population of workers.

3.2 Aggregate trends in outsourcing

We start by asking by how much has outsourcing risen in France. Figure 2(a) shows that outsourcing expenditures as a fraction of aggregate payroll 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 subgategory of outsourcing that we can reliably measure past 2008: temporary workers. Of course, extrapolating the trend in overall outsourcing expenditures using the 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(a) reveals that outsourcing may account between 10% and 20% of the aggregate wage bill in France by 2016.[14]

We complement this expenditure-based measure with employment-based metrics because outsourcing expenditures do not distinguish between spending on low or high skill outsourced workers and conflate prices and quantities. We follow Goldschmidt and Schmieder (2017) and rely on industry and occupation codes to detect contractor firms and service workers at contractor firms. We define low skill contractors as firms that operate in 3-digit industries who provide one of the following services to other firms: security, cleaning, food, interim and general administrative services, and call centers. We consider firms that provide accounting, law or consulting services for high skill contractors. Similarly, we define low and high skill service workers as workers in occupations related to these same services.[15]

We find that rising outsourced employment accompanies growing outsourcing expenditures. Figure 2(b) 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, and is substantially more pronounced than for high skill workers. Figure 2(c) also reveals that the increase in the employment outsourcing share is driven specifically by service workers reallocating toward contractor firms over time. The fraction of low skill

[13] France 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.

[14] That 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.

[15] In the *Nomenclature Française d'Activité Rev. 1*, these industries correspond to codes 74.1, 74.5, 74.6, 74.7, 74.8 and 55.5. These occupations correspond to codes 377a, 468a, 488a, 561b, 561c, 561d, 636d, 686a, 686b, 534a, 534b, 461b, 461c, 461f, 542a, 542d, 543g, 543h, 541d, 461d, 461e, 543b, 543c, 543e, 372e.

Figure 2: Outsourcing share over time in France

(a) Expenditure share. (b) Employment at contractors. (c) Service emp. at contractors.

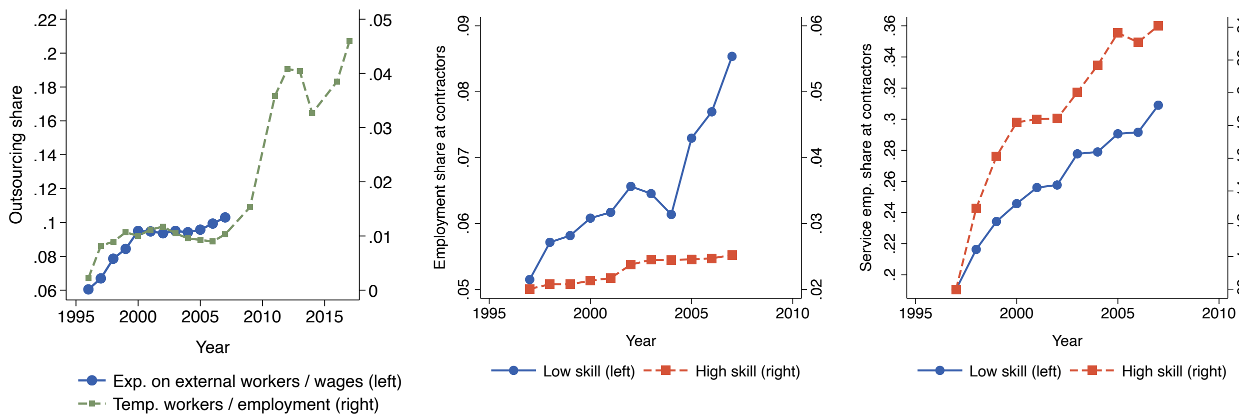

Note: Panel 2(a): firm-level dataset (FICUS-EAE-DADS *Postes*-Trade). Solid blue line: aggregate outsourcing share, computed as aggregate expenditures on external workers relative to the aggregate in-house wage bill. Dashed green line: temporary work share, computed as total temporary workers at firms with more than 100 in-house employees, relative to all employment at firms with more than 100 in-house employees. Panels 2(b) and 2(c): worker-level dataset (DADS *Panel*). Panel 2(b): fraction of all employment at low skill contractors and high skill contractors. Panel 2(c): fraction of service employment at low skill contractors and high skill contractors. 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.

service workers employed at contractors rises by 16 p.p. over the decade we study.[16]

Table 9 in Online Appendix F.1 describes salient features of firms that rely on outsourcing to firms that do not. Consistent with Propositions 3 and 4, firms that outsource are larger, sell more and have higher value added than firms that do not outsource. 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. By contrast, Tranportation 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 predictions of our theory related to the distributional effect of outsourcing. To that end, we use the worker-level dataset (DADS *Panel*). Following our results in Section 2.3, we test whether contractors locate at the bottom of the job ladder by paying lower wages, hiring less from employment, having higher churn and lower net poaching.[17]

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),

[16]Our 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.

[17]We also check that, consistent with an upward-sloping labor supply curve, there is a size-wage premium for service workers in Online Appendix F.3. Importantly, 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.

Figure 3: Contractors and non-contractors on the job ladder.

(a) Raw wages.

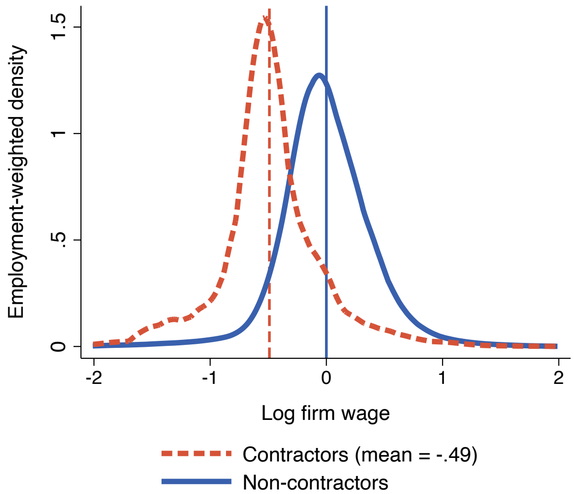

(b) Firm premia.

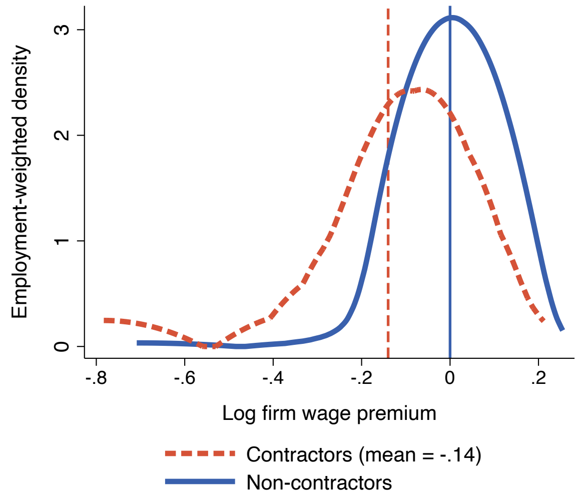

(c) Frac. hires from emp.

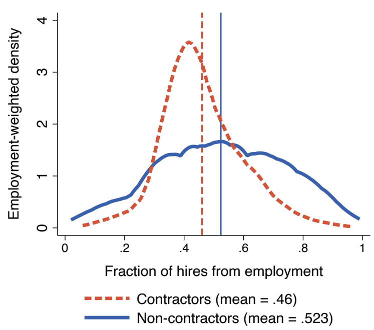

(d) Churn.

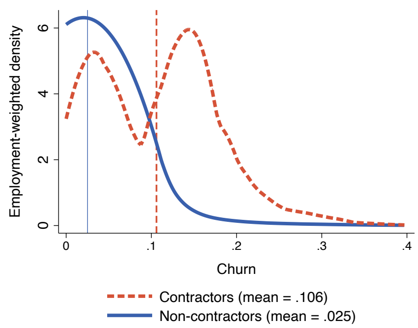

(e) Net poaching.

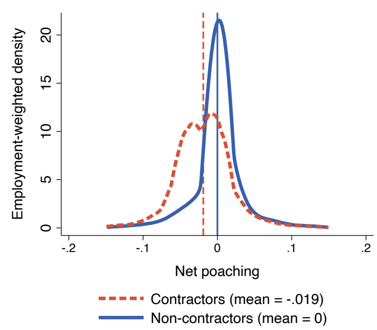

Note: Kernel density plots of labor market statistics of contractor firms (dashed orange) vs. non-contractor firms (solid blue) in worker-level dataset (DADS *Panel*). Vertical lines depict means. Raw wages: firm-level mean wages. Firm premia: firm fixed effect from the AKM regression (9). Fraction hires from employment: hires from employment divided by total hires as in equation (5). Churn: total separations divided by employment as in (6). Net poaching: hires from employment net of quits to employment divided by total employment as in (7).

henceforth AKM:

$$\log w_{i,t} = \varphi_i + \psi_{J(i,t)} + \eta_{i,t}, \tag{9}$$

where $\psi_{J(i,t)}$ is a firm fixed effect. Workers who move between firms identify separately worker and firm fixed effects if worker mobility is conditionally random (Card, Heining, 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 unconditional mean worker and mean firm wages. We then estimate equation (9) 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 3(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 3(b). The penalty drops to 14% on average after removing worker fixed effects with specification (9). 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. (2023) 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 toward 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 3(c). The fraction of hires from employment is 6 p.p. lower at contractors. Figure 3(d) indicates that churn at contractors is 8 p.p. above that at other firms. Finally, contractors tend to lose their workers to other firms as shown in Figure 3(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 confirmed 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. Proposition 3 indicates that when productivity z rises conditional on outsourcing costs $\boldsymbol{\varepsilon} = \{\varepsilon_s\}_s$, firms spend relatively more on outsourcing: $\mathcal{E}^*(z, \boldsymbol{\varepsilon})$ increases.

Consistent with Proposition 3, we use expenditures on external workers $\mathcal{E}_{ft}$ for firm f in time period t as our first measure of outsourcing. Since there may be reasons beyond our model that drive a mechanical correlation between firm scale and outsourcing expenditures, we also use the outsourcing share S_{ft} as an alternative dependent variable. We define it as 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: $S_{ft} = \frac{\mathcal{E}_{ft}}{\mathcal{W}_{ft}+\mathcal{E}_{ft}}$. Using the outsourcing share is a more demanding test of our theory because our model does not guarantee that the outsourcing share should always be increasing in productivity.[18]

In practice, we do not observe revenue productivity z directly, but we may observe variables that correlate with it—for instance, export demand shocks. In our framework, revenues $R^*(z, \boldsymbol{\varepsilon})$ are increasing in productivity conditional on outsourcing costs $\boldsymbol{\varepsilon}$. If we can isolate variation in revenue productivity z that is unrelated to outsourcing costs $\boldsymbol{\varepsilon}$, we can thus use revenues as a relevant scale to describe the selection effect.

To construct the relevant empirical counterpart of revenues, we use measures of value added plus outsourcing expenditures given that our framework does not feature other intermediate inputs: revenues of firm f in time period t are $R_{ft} = \mathrm{VA}_{ft} + \mathcal{E}_{ft}$, where VA_{ft} is value added and $\mathcal{E}_{ft}$ is expenditures on external workers. While these definitions are those that are consistent with our theory, we replicate

[18]For 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 14 in Online Appendix F.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 4: Outsourcing by revenues.

(a) Outsourcing expenditures. (b) Outsourcing share.

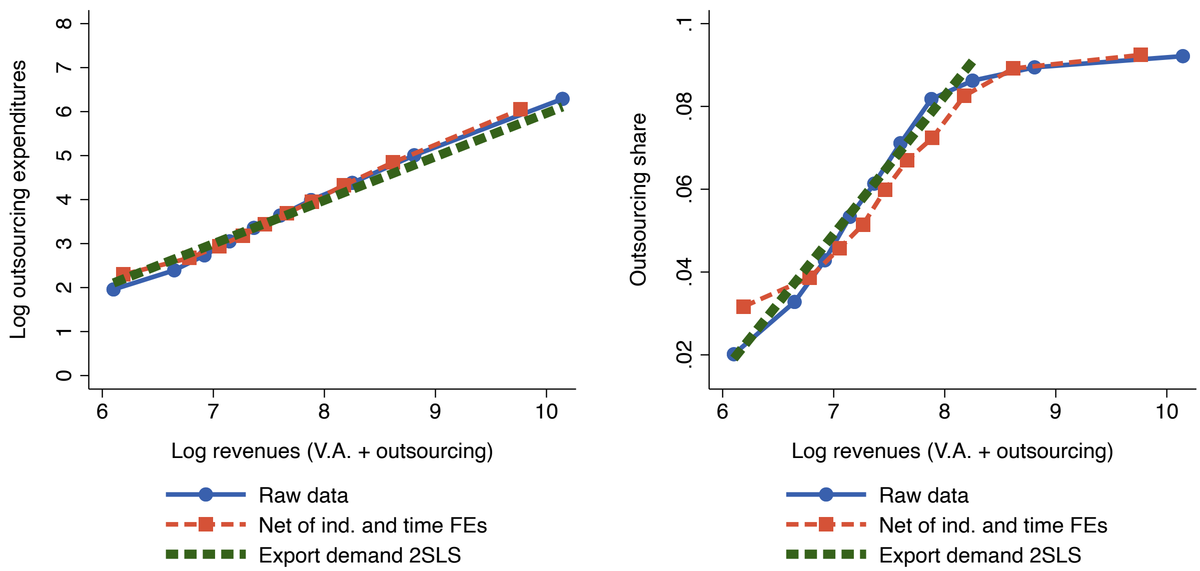

Note: Solid blue line: raw data. Dashed orange line: after removing 3-digit industry and time period fixed effects from log outsourcing expenditures, the outsourcing share and log value added. Thick dashed green line: 2SLS estimate using the export demand shift-share instrument in equation (11). Panel (a): log outsourcing expenditures. Panel (b): outsourcing share. Firm-level dataset (FICUS-EAE-DADS *Postes*-Trade).

all our results using value added VA_{ft} instead of revenues in Appendix B and Online Appendix F.1 to ensure that any mechanical correlation in the definition of our variables is not driving our results, and find virtually unchanged results.

With these metrics in hand, we investigate whether firms select into outsourcing by testing if $\frac{\partial \log \mathcal{E}(z,\boldsymbol{\varepsilon})}{\partial z} \Big/ \frac{\partial \log R(z,\boldsymbol{\varepsilon})}{\partial z} > 0$ and $\frac{\partial S(z,\boldsymbol{\varepsilon})}{\partial z} \Big/ \frac{\partial \log R(z,\boldsymbol{\varepsilon})}{\partial z} > 0$. Identification crucially depends on the co-movement between productivity z and outsourcing costs $\boldsymbol{\varepsilon}$ across firms.

In the cross-section one might expect outsourcing costs to be correlated with productivity—for instance because better management makes firms more productive but is also better at leveraging contractors. We address these endogeneity concerns in two steps. First, we control for time-invariant unobserved differences in outsourcing costs in the two-way fixed effect regression:

$$Y_{ft} = \alpha_t + \beta_f + \gamma \log R_{ft} + \eta_{ft}, \qquad Y_{ft} \in \{\log \mathcal{E}_{ft}, S_{ft}\}, \tag{10}$$

where α_t is a time period fixed effect, β_f a firm fixed effect, and η_{ft} a mean zero residual.

Second, we turn to an instrumental variable strategy to address the possibility that outsourcing costs co-move with productivity over time within firms. We use shocks to revenue productivity z that are plausibly unrelated to outsourcing cost shocks $\boldsymbol{\varepsilon}$. We construct foreign export demand shocks by exploiting the granularity of our customs data and following 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

Figure 5: Selection into outsourcing: Instrumental variable approach.

(a) First stage. (b) Expenditures: 2SLS. (c) Share: 2SLS.

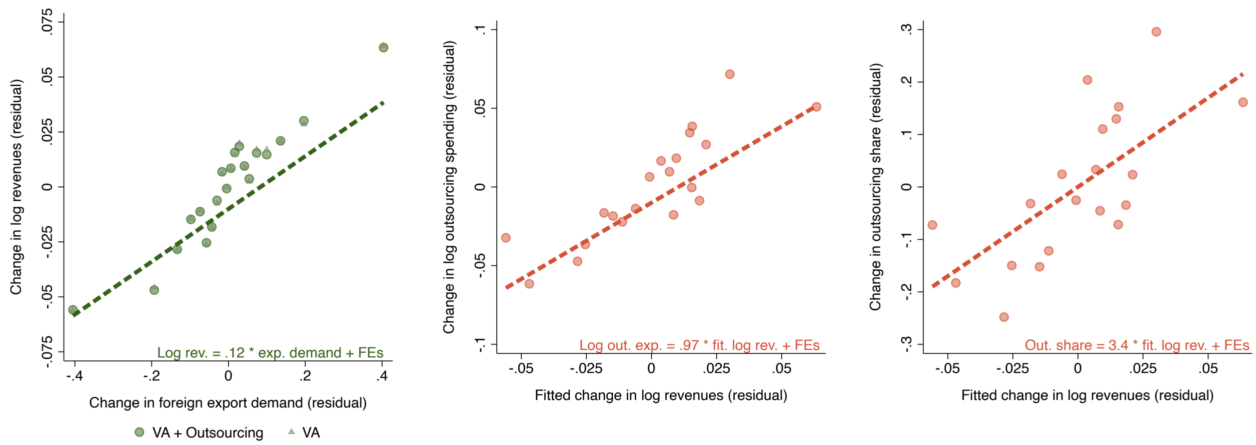

Note: Bin-scatterplot of the first stage (panel (a)) and two stage least square estimates for selection into outsourcing with log outsourcing expenditures (panel (b)) and the outsourcing share (panel (c)) as dependent variable. Figure 13 in Online Appendix F.1 depicts the reduced-form relationship between outsourcing and export demand. Panels include regression lines for 20 bins. Coefficients may differ from full sample regression coefficients reported in Tables 3 and 4, Appendix B. Firm-level dataset (FICUS-EAE-DADS *Postes*-Trade).

as a shift-share:

$$Z_{f,t} = \sum_j \pi_{f,t_0,j} \, \Delta \log X_{j,t,-f}. \tag{11}$$

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 changes in 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 12, Online Appendix F.1. Up to a first-order approximation, the 2SLS estimate of γ in equation (10) then captures the ratio of average partial derivatives $\mathbb{E}\left[\frac{\partial \log \mathcal{E}(z,\varepsilon)}{\partial z}\right] / \mathbb{E}\left[\frac{\partial \log R(z,\varepsilon)}{\partial z}\right] > 0$ and $\mathbb{E}\left[\frac{\partial S(z,\varepsilon)}{\partial z}\right] / \mathbb{E}\left[\frac{\partial \log R(z,\varepsilon)}{\partial z}\right] > 0$, and thus selection into outsourcing.

Figure 4 descriptively plots outsourcing by decile of revenues in our firm-level dataset (FICUS-EAE-DADS *Postes*-Trade). Figure 4(a) reveals that in the raw data, high revenue firms have higher outsourcing expenditures. The relationship is nearly log-linear for outsourcing expenditures. Figure 4(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. Likewise, the relationship between outsourcing expenditures or the outsourcing share on the one hand, and revenues on the other hand, is largely unchanged when we focus on within-firm variation in Tables 3 and 4, Appendix B.

Figure 5 represents our instrumental variable strategy graphically. Panel (a) shows a strong first stage: export demand growth predicts revenue growth, with an F-statistic of 332 well above conventional thresholds for weak instruments (Stock and Yogo, 2005). Panels (b) and (c) uncover a positive relationship between revenue productivity and outsourcing driven by export demand shocks.

We collect our estimates using our firm-level sample in Tables 3 and 4, Appendix B. A productivity shock driven by export demand that implies a 10% increase in revenues leads outsourcing expenditures to rise by 10% and the outsourcing share to rise by 0.33 p.p.[19] All our estimates are economically and statistically significant at the 0.1% level. The thick dashed green lines in Figure 4 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 11 in Online Appendix F.1 together with column (6) of Tables 3 and 4 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 3 and 4 in Appendix B 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.[20]

3.5 The productivity effect

We now test the third core prediction of our theory: the productivity effect of outsourcing. Proposition 4 indicates that when outsourcing costs ε_s decline so much that firms outsource, they expand: $R(z, \boldsymbol{\varepsilon})$ rises. Consistent with our framework, we use revenues R_{ft} as our dependent variable in the main text. To ensure that our results are not mechanically driven by a rise in outsourcing expenditures, we also use value added VA_{ft} as an alternative dependent variable in Appendix B and Online Appendix F.2 and find virtually identical results. The independent variable is the outsourcing share S_{ft}. We seek to identify whether $\frac{\partial \log R(z,\boldsymbol{\varepsilon})}{\partial(1/\varepsilon_s)} > 0$. Since we do not measure outsourcing costs directly, we ask if $\frac{\partial \log R(z,\boldsymbol{\varepsilon})}{\partial(1/\varepsilon_s)} \Big/ \frac{\partial S(z,\boldsymbol{\varepsilon})}{\partial(1/\varepsilon_s)} > 0$.

The key identification challenge is to isolate changes in outsourcing expenditures driven by outsourcing costs rather than revenue productivity. As in Section 3.4, we implement an instrumental variable strategy to address the possible correlation between outsourcing costs and productivity. We leverage that firms are differentially exposed to service occupations o: food, security, cleaning or general administrative occupations. Our first step is to construct average outsourcing expenditures on occupation o, denoted by $\Phi_{o,t,-f}$. We interact initial payroll shares ϕ_{f,o,t_0} with firm-level outsourcing

[19]Since our instrument only affects exporters, we also confirm that exporters exhibit a similar within-firm OLS relationship between revenues and outsourcing relative to all firms in column (5) in Tables 3 and 4, Online Appendix F.1. In panel 5(c), the coefficient is 0.34 p.p. when we run the regression at the bin level rather than at the firm level.

[20]Abraham 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 18(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 216,051 firm-periods.

Figure 6: Productivity effect: Instrumental variable approach.

(a) First stage. (b) Reduced form. (c) 2SLS.

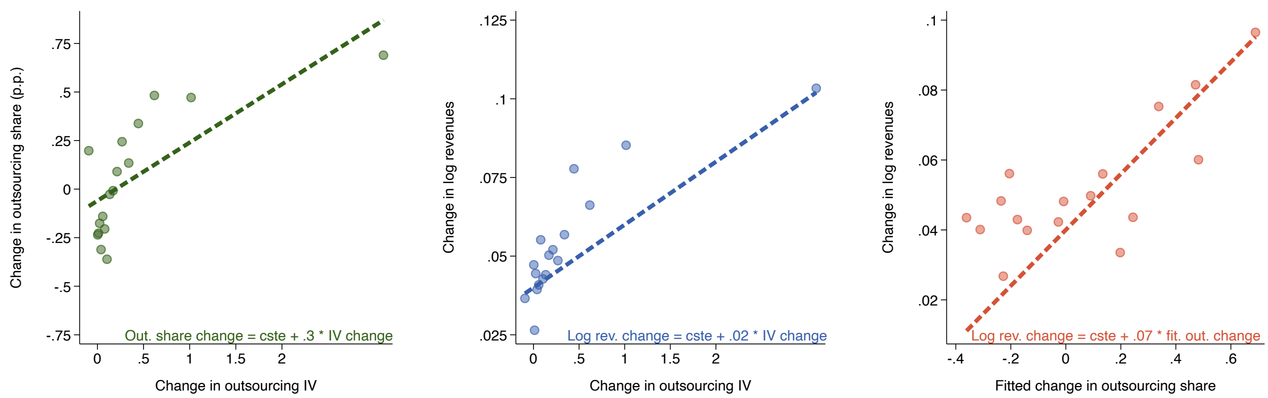

Note: Bin-scatterplot of the first stage (panel (a)), reduced form (panel (b)) and two stage least square (panel (c)) estimates for the productivity effect using revenues R_{ft} as a dependent variable. Panels include regression lines for 20 bins. Coefficients may differ from full sample regression coefficients reported in Table 5, Appendix B. Results using value added VA_{ft} as a dependent variable in Figure 16, Online Appendix F.2. Firm-level dataset (FICUS-EAE-DADS *Postes*-Trade).

expenditures $\mathcal{E}_{f,t}$. We then sum across firms to obtain $\Phi_{o,t,-f} = \frac{1}{N_{-f}} \sum_{f' \neq f} \phi_{f',o,t_0} \mathcal{E}_{f',t}$. 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 $\Phi_{o,t,-f}$: $\widehat{S}_{f,t} = \frac{\sum_o \phi_{f,o,t_0} \Phi_{o,t,-f}}{\mathcal{W}_{f,t_0} + \mathcal{E}_{f,t_0}}$.

We then consider the linear regression framework:

$$\log R_{ft} = \alpha'_t + \beta'_f + \gamma' S_{ft} + \eta'_{ft}. \tag{12}$$

We control for time-invariant differences in productivity with a firm fixed effect. We instrument changes in the outsourcing share $S_{f,t}$ with changes in the predicted outsourcing share $\widehat{S}_{f,t}$ to address the possible correlation between time-varying outsourcing costs and productivity.

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 ϕ_{f,o,t_0}, changes in outsourcing expenditures $\Delta\Phi_{o,t,-f}$ and the resulting instrument $Z'_{f,t}$ in Figure 15 in Online Appendix F.2. Up to a first-order approximation, the 2SLS estimate of γ' in (12) identifies the ratio of average partial derivatives $\mathbb{E}\left[\frac{\partial R(z,\boldsymbol{\varepsilon})}{\partial(1/\varepsilon_s)}\right] / \mathbb{E}\left[\frac{\partial S(z,\boldsymbol{\varepsilon})}{\partial(1/\varepsilon_s)}\right]$ and thus the productivity effect of outsourcing.

Figure 6 represents our instrumental variable strategy graphically using our firm-level dataset (FICUS-EAE-DADS *Postes*-Trade). Panel 6(a) reveals a strong and positive first stage. Panel 6(b) reveals that growth in the predicted outsourcing share also leads to growth in firm-level value added. As a result, panel 6(c) uncovers a positive relationship between outsourcing and revenues.

We collect our estimates using our full firm-level sample in Table 5, Appendix B. The first stage F-statistic is 23.8. A decline in outsourcing costs that implies a 1 p.p increase in the outsourcing share leads to 9% growth in revenues. This point estimate is economically and statistically significant at the 0.1% level. We conclude that declines in outsourcing costs have positive productivity effect at the firm level: firms that outsource produce more.

3.6 Alternative explanations

In principle, mechanisms that our theory does not emphasize may also lead firms to outsource. In this subsection, we examine whether three prominent alternative explanations may be key drivers of outsourcing in France.

Volatility—Firms may outsource because they value the associated flexibility when demand is volatile. We find some evidence in favor of a modest relationship between volatility and outsourcing, consistent with Abraham and Taylor (1996). Table 11 in Online Appendix F.4 suggests that firm-level value added volatility is positively associated with the outsourcing share, though the standardized coefficient is two to three times smaller than for firm scale. Thus, firm scale appears to be a stronger determinant of outsourcing than workforce flexibility.[21]

Equity—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 is partly generated by equity concerns, but we still investigate whether equity concerns per se lead to outsourcing. Under equity motives, firms with more unequal pay structures have the strongest incentives to outsource. Table 12 in Online Appendix F.4 indicates, if anything, the opposite: firms with more unequal pay structures outsource less, not more. Therefore, equity concerns do not appear to affect outsourcing more than firm scale.[22]

Unions and size-based regulations—The upward-sloping labor supply curve could also be the result of union wage-setting or size-based labor market regulations. While our theory is equally valid in this case, size-based regulations are 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.[23] We use a regression discontinuity design to assess the role of unions and size-based regulations on outsourcing. Figure 18(a) in Online Appendix F.4 shows that there is no statistically significant nor economically meaningful discontinuity in the outsourcing share around the 50 employees

[21] At a more aggregated level, Figure 17 in Online Appendix F.4 shows that, if anything, small firms or industries that are more volatile outsource less, not more.

[22] Our results are consistent with the growing literature that documents equity concerns (Card, Mas, et al., 2012; Breza et al., 2017), which finds that equity concerns are primarily binding across workers within the same occupation, rather than across occupations or worker performance categories.

[23] They 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"*).

threshold, despite French firms indeed bunching at that cutoff. Hence, union wage-setting or size-based regulations do not seem to strongly affect outsourcing behavior in France.[24]

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

4 Extended model and estimation

To structurally evaluate the impact of rising outsourcing on the economy, we enrich our environment along several dimensions before discussing our estimation strategy.

4.1 Quantitative setup

Extensions—We expand our framework in three ways. First, to capture the idea that goods producers may expand their human resource departments to hire more in-house without increasing wages, we let goods producers create any number of jobs—or post any number of vacancies—v in each market s. This job creation effort comes at a convex cost $c_s(v) = \frac{c_{0s}}{1+\gamma} v^{1+\gamma}$, for $\gamma > 0$. When $\gamma \to +\infty$ we recover the model of Section 2. The number of workers a goods producer attracts and retains now reflects its vacancy share:

$$n_s(w, v) = \frac{(1 + k_s)e_s}{\big(1 + k_s(1 - F_s(w))\big)^2} \cdot \frac{v}{V_s}, \tag{13}$$

where the equilibrium number of vacancies in market s satisfies: $V_s = \mathcal{V}_s + M^g \sum_{\boldsymbol{o}:o_s=0} \int v_s(z, \boldsymbol{\varepsilon}, \boldsymbol{o})\Gamma(dz, d\boldsymbol{\varepsilon})$. $\mathcal{V}_s$ is the measure of vacancies posted by contractor firms. $v_s(z, \boldsymbol{\varepsilon}, \boldsymbol{o})$ denotes the number of vacancies posted by a firm with productivity z, outsourcing costs $\boldsymbol{\varepsilon}$ and decision $\boldsymbol{o} = \{o_s\}_s \in \{0, 1\}^S$. With a discrete set of skills, goods producers then solve:

$$\begin{aligned} \pi(z, \{\varepsilon_s\}_s) \quad &= \quad \max_{\substack{\{n_s\}_s, \{v_s\}_s, \\ \{w_s\}_s, \{o_s\}_s}} R\big(z, \{n_s\}_s\big) - \sum_{s=1}^{S} \Big\{ \big[(1 - o_s)w_s + o_s p_s \varepsilon_s\big] n_s + (1 - o_s) c_s(v_s) \Big\} \\ &\text{s.t.} \quad n_s = n_s(w_s, v_s) \text{ as per (13) if } o_s = 0. \end{aligned}$$

Second, we let contractors differ in productivity, which allows us to match their wage distribution exactly. In addition, to capture the idea that contractors may specialize recruiting activities, we endow them with a possible comparative advantage in hiring. Their recruiting cost function is $c^c(v) = \frac{c_{0s}}{1+\gamma} \bar{c}^{\gamma} v^{1+\gamma}$. The relative marginal cost $\bar{c} \lesseqgtr 1$ 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.

Third, we endogenize the mass of producers and contractors through a free entry condition. As in Melitz (2003), firms pay an entry cost before knowing their productivity. The free entry conditions $\mathbb{E}[\pi(z, \{\varepsilon_s\}_s] = \eta$ and $\mathbb{E}[\pi^c(z)] = \eta^c$ determine the equilibrium mass of producers and contractors, where η and η^c are their respective entry costs.

[24] This 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.

Parametric choices—We impose the following parametric assumptions. We solve the model for three skill types $S = 3$. We use a Cobb-Douglas revenue function nested in a decreasing returns upper tier: $R\big(z, \{n_s\}_s\big) = \Big(z \prod_{s=1}^{S} n_s^{a_s}\Big)^{\rho}$, $\sum_{s=1}^{S} a_s = 1$. 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 ($s = 1$).

The joint distribution of goods producer productivity and outsourcing costs (z, ε_1) 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_1 as an iceberg trade cost. The distribution of contractor productivity is log-normal with mean zero and standard deviation ν^c.

We specify a Cobb-Douglas matching function: $\mathcal{M}_s = \mu_s(m_s(u_s + \zeta_s(1-u_s))^{\xi} V_s^{1-\xi}$. 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. We relegate additional derivations to Online Appendix G.1 and computation details to Online Appendix G.3.

4.2 Identification

To estimate our quantitative model transparently, we prove an identification result. We extend the non-parametric identification results of Bontemps et al. (1999) under constant returns and perfect substitutability to our setting with decreasing returns to scale, complementarities in production and outsourcing.

Proposition 6. *(Identification)*
Given the matching function elasticity ξ, skill-level worker flow data, firm-level value added and outsourcing expenditures data, and firm-by-skill employment and wage data, the parameters, $\{\delta_s\}_s$, $\{\mu_s\}$, $\{\zeta_s\}_s$, $\{m_s\}_s$, $\{a_s\}_s$, ρ, γ, ν^c, $\bar{c}$, η, and η^c are identified.

Proof. See Online Appendix H.2. □

Proposition 6 identifies 21 out of the 25 parameters of the model given the matching function elasticity ξ. Its proof is recursive and constructive. The first step is independent of the production function and is identical to Bontemps et al. (1999). Labor market transitions identify the parameters governing search frictions. The 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, which identifies the matching function efficiency μ_s. The employment-to-employment transition rate EE_s identifies the arrival rate $\lambda_s^E = k_s \delta_s$ after accounting for rejected offers along rungs of the job ladder: $\frac{\mathrm{EE}_s}{\mathrm{EN}_s} = (1 + 1/k_s)\log(1 + k_s) - 1$. Given search frictions, the employment share of skill s identifies the mass of workers m_s. Finally, the replacement rate determines unemployment insurance b_s.

The second step adapts the insights of Bontemps et al. (1999) to our setting. Specifically, we invert the first order condition behind the wage equation in Proposition 1 given data on wages w_{js} and firm ranks F_{js} to recover the equilibrium distribution of marginal products of labor, which we denote MPL:

$$\text{MPL}_{js} = w_{js} + \left(\frac{1 + k_s(1 - F_{js})}{2k_s}\right)\frac{\partial w_s(F_{js})}{\partial F}, \tag{14}$$

where j denotes a firm, F_{js} is its rank in the wage offer distribution from skill s, and $w_s(F)$ is the inverse function of $F_s(w)$. The second term in equation (14) captures the markdown: the wedge between the marginal product of labor and wages due to search frictions introduce.

Equation (14) identifies the marginal product of labor for any production function and all firms j. Our approach then departs from Bontemps et al. (1999). We use the definition of the marginal product of labor: $\frac{\text{MPL}_{js}}{R_j/n_{js}} = \rho a_s$ for in-house producers, and $\frac{\text{MPL}_{js}}{R_j/n_{js}} = \frac{\rho a_s}{1-\rho a_1}$ for outsourced producers. Using data on employment by firm and skill and value added by firm, taking appropriately weighted means across firms identifies ρ and $\{a_s\}_s$.

The combination of firm size and wage rank provides an estimate of the number of vacancies each firm posts in equilibrium, v_{js}, as per equation (13). Using these estimates, the vacancy optimality condition for producers, $\log v_{js} = \overline{\alpha}_s + \frac{1}{\gamma}\log\left[n_{js}\left(\text{MPL}_{js} - w_{js}\right)\right]$, identifies the vacancy cost function elasticity γ by matching the within-skill across-firm variation in estimated vacancies and estimated returns to vacancy creation. Given the estimation for the elasticity γ, the vacancy optimality condition for contractors then determines the relative vacancy cost of contractors $\overline{c}$.

With these parameters in hand, we construct equilibrium aggregate outsourcing expenditures and profits using our estimates of equilibrium objects without having to solve the model. Doing so is feasible given observed choices only—value added, employment and outsourcing—without having to estimate the parameters of the producer productivity distribution. Average firm size and average firm employment allow us to recover the equilibrium total mass of firms in the economy $M^g + M^c$ after accounting for selection into entry. Aggregate outsourcing expenditures via the outsourcing market clearing condition (8) lets us recover the equilibrium relative mass of contractors M^c/M^g. Given all the previous parameters, average profits of producers and contractors identify entry costs.

Given technology parameters and firm-level data on value added, we invert the production function to back out productivity z_j for all firms, producers and contractors. Equipped with productivity estimates and given constant returns for contractors, we directly identify the standard deviation of contractor productivity ν^c. For producers, the additional selection into outsourcing complicates the mapping to the productivity distribution. We relegate the estimation of its parameters to a separate estimation step in Section 4.3, where we also describe how we ready the data.

4.3 Estimation strategy

We set a quarterly frequency. We define low ($s = 1, 2$) and high ($s = 3$) skill worker groups based on their primary non-farm 1-digit occupation code. High skill workers are workers in business, management and intermediate jobs (occupational classification codes 2,3 and 4). Low skill workers are other employees and manual workers (occupation classification codes 5 and 6). Within low-skill workers, we define "service" workers ($s = 1$) as workers who are likely to be exposed to an outsourcing event throughout their career. In practice, we define this category as workers who have ever worked in the service occupations as defined in Section 3.2, as well as workers who have ever worked in other occupations are reached by at least one direct job-to-job transition from service occupations.[25] Other low skill workers are "core" workers. 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 group firms in the data into three categories: contractors as identified by their 3-digit industry code, producers who outsource, and producers who do not outsource.

With these definitions in hand, we start by setting the elasticity of the matching function $\xi = 0.5$ to a value commonly used in the literature since our data does not let us estimate it credibly (Petrongolo and Pissarides, 2001). We estimate the next 21 parameters using Proposition 6. We estimate the remaining 4 parameters—τ, ν, σ, and ι—jointly with a Method of Simulated Moments (MSM) estimator. We propose a heuristic argument that supports identification.

Absent selection into outsourcing, the dispersion of revenue productivity for in-house and outsourced producers would directly map into the unconditional dispersion of productivity and outsourcing cost, ν and σ. We therefore calibrate ν and σ to match these two moments obtained during the model inversion after accounting for selection in the model.

The correlation between productivity and outsourcing costs informs the relationship between outsourcing and value added: the larger the correlation, the more likely is an outsourcing firm to have low productivity and therefore low value added. Accordingly, we calibrate ι to match the cross-sectional OLS relationship between the outsourcing share and valued added.[26]

Finally, the outsourcing cost τ determines the productivity wedge between contractors and producers. Since wages track the marginal product of labor, we set τ to match the outsourcing wage penalty for low skill service workers. We provide additional details in Online Appendix H.3.

Figure 21 in Online Appendix H.4 confirms our identification argument numerically. Given estimates of ν, σ and ι, we can exactly invert the model to recover τ (see Online Appendix H.3) and thus only need to establish numerical identification for ν, σ and ι. Each of the three parameters affects the particular moment that we use to identify it in isolation, as well as increases the joint objective function away from its minimum.

4.4 Estimation results

Table 1 summarizes our parameter estimates. We obtain standard values for standard parameters. The revenue function curvature parameter is $\rho = 0.77$, and the curvature of the vacancy cost is $\gamma = 1.68$. Both values are well within the ranges found in the literature. 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, consistent with the literature on the skill premium.

The estimates of parameters linked to outsourcing point to comparative advantage of contractors grounded in job creation rather than in production of labor services. 21% of the total workforce ($m_1 = 0.21$) is subject to outsourcing, which represents 42% of low-skill workers. We estimate that

[25] We also restrict these other occupations to be large enough and constitute at least than 1% of the workforce.

[26] We target the cross-sectional relationship for two reasons. First, within-firm estimates are short-term elasticities whereas our model is in a long-term steady state. Second, our 2SLS estimates are local average treatment effects that are more challenging to relate to the model. Nevertheless, we compare our estimated model to our 2SLS estimates in Section 4.4.

Table 1: Parameter estimates and empirical targets.

Parameter	Interpretation	Target	Empirical moment	Simulated moment	Parameter estimate
		A. External calibration			
χ	Matching function elasticity				0.50
		B. Model inversion			
$\{\delta_s\}_s$	Job loss rate	EN rate	$\{0.02, 0.01, 0.02\}$		$\{0.02, 0.01, 0.02\}$
$\{\mu_s\}_s$	Matching efficiency	NE rate	$\{0.84, 0.82, 1.50\}$		$\{0.20, 0.13, 0.16\}$
$\{\zeta_s\}_s$	Relative search efficiency	EE rate	$\{0.87, 0.48, 0.85\}$		$\{0.18, 0.06, 0.14\}$
$\{m_s\}_s$	Mass of workers	Employment share	$\{0.21, 0.29, 0.51\}$		$\{0.21, 0.28, 0.51\}$
$\{b_s\}_s$	Unemployment insurance	Replacement rate	$\{0.40, 0.40, 0.40\}$		$\{0.40, 0.30, 0.46\}$
$\{a_s\}_s$	Skill productivity	Skill average MPL	$\{0.22, 0.21, 0.57\}$		$\{0.22, 0.21, 0.57\}$
ρ	Decreasing returns to scale	Average MPL	0.77		0.77
γ	Elasticity vac. cost	Equation (61)	1.68		1.68
ν^c	St. dev. contractor TFP	St. dev. contractor MPL	0.27		0.27
$\bar{c}$	Contractor vacancy cost	Equation (62)	0.06		0.06
η	Producer entry cost	Producer avg. profits	42.8		42.8
η^c	Contractor entry cost	Contractor avg. profits	2.23		2.23
		C. Indirect inference			
τ	Outsourcing cost	Outsourcing wage penalty	0.88	0.88	0.42
ν	St. dev. producer TFP	St. dev. in-house TFP	0.41	0.41	0.41
σ	St. dev. out. costs	St. dev. outsourced TFP	0.36	0.37	0.33
ι	Covariance btw. z and ε	Revenues-out. share relationship	1.73	1.74	0.01

Note: Firms in the data grouped into three categories: contractors as identified by their 3-digit industry code, producers who outsource, and producers who do not outsource. Within each category, firms grouped into 100 bins of average unconditional wage to reduce measurement error and project the data on the relevant dimension of heterogeneity. Grouping consistent with the model in which more productive firms pay higher wages. Outsourcing wage penalty: 12% instead of 14% due to grouping and sample selection. Additional details in Online Appendix H.1.

contractors are less efficient than goods producers at producing low-skill labor services: the estimated comparative advantage parameter $\tau = 0.42$ implies that the employment-weighted average productivity of contractors is 50% lower than that of producers for low-skill labor services. Thus, given our framework in which wages imperfectly track the marginal product of labor, the observed outsourcing wage penalty implies a much larger discrepancy in productivity terms.

By contrast, we estimate contractors to be substantially more efficient than producers at job creation: their relative cost shifter is 0.06, implying that their marginal cost of job creation is 12% of the producers' in equilibrium.[27] Overall, the combination of large contractor size and a positive wage penalty in the data favors an interpretation in which contractors are specialists in recruiting activities rather than in production activities. We interpret these findings as evidence in favor of the cost-saving view rather than the comparative advantage view.

To lend further support to this interpretation, we gauge the fit of the model and its ability to match non-targeted moments in Table 2. We first examine wage dispersion and associated labor

[27]The 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.

Table 2: Model fit and over-identification statistics

	Data			Model		
	A. Economy-wide					
St. dev. wage		0.27			0.24	
Between skill (%)		69			59	
Within skill (%)		31			41	
Selection IV		3.34	(1.57, 5.11)		1.93	
Productivity IV		0.08	(0.04, 0.12)		0.02	
Labor supply elasticity		4.3—6.5			4.8	
	B. By producer type					
	Contractors	Producers	Diff.	Contractors	Producers	Diff.
Firm premia	-0.12	0.00	-0.12	-0.12	0.00	-0.12
Frac. hires from emp.	0.46	0.52	-0.06	0.65	0.71	-0.05
Churn	0.11	0.03	0.08	0.07	0.04	0.03
Net poaching	-0.02	0.00	-0.02	-0.02	0.00	-0.02

Note: range of point estimates of labor supply elasticities in the data from Lamadon et al. (2022) for the U.S. Labor supply elasticity in the model: firm-level average of $wn'_s(w)/n_s(w)$ for $s = 1$. Weighted average across all skills: 4.5. Selection and productivity: point estimate and 95% confidence interval.

market moments in panel A. Our estimation initially uses wage distributions to estimate marginal products of labor, but then restricts the freedom of the model by imposing a common productivity by firm and a common revenue function across firms. Despite these restrictions, we show that the model matches the overall standard deviation of wages in the data reasonably well, as well as its split within and between skills.

The estimated job ladder model connects the targeted outsourcing wage penalty and relative contractor size—through comparative advantage in production τ and job creation $\bar{c}$—to untargeted differences in labor market reallocation between producers and contractors through Corollary 1. We examine these moments in panel B. The model predicts somewhat more hiring from employment than in the data across the board, but the gap between contractors and producers in the model (-0.05) is close to its counterpart in the data (-0.06). Relative net poaching by contractors is virtually identical in the model and in the data at -0.02, while relative churn is positive but lower in the model (0.03) than in the data (0.08).

Finally, we examine two firm outcomes specifically tied to the decision to outsource in panel A: selection into outsourcing and the productivity effect. We target the cross-sectional OLS relationship between the outsourcing share and firm value added in the data. When contrasting within-firm estimates with the model, we must take a stance on the incidence of shocks. We choose the size of the shock to match the within-firm OLS selection coefficient by considering a joint change in productivity z and outsourcing costs ε. We then ask whether the IV coefficients align with the data by considering only a change in z for selection into outsourcing or only a change in ε for the productivity effect. We provide more details in Online Appendix H.5. We report the average effect across firms in Table 2 and the heterogeneous treatment effects in Figure 22, Online Appendix H.5.

The covariance ι between productivity z and outsourcing costs ε controls the gap between OLS and IV estimates for selection into outsourcing. Consistent with our positive but moderate estimate of the covariance $\iota = 0.01$, our IV estimate at 1.93 is larger than the OLS estimate as in the data. Although it remains below the point estimate in the data at 3.34, it is not statistically different at the 5% level given our 95% confidence interval (1.57, 5.11). We conclude that the fundamental drivers of selection into outsourcing—bypassing high labor costs—suffice to align the model with the data without the need for a large covariance ι.

Decreasing returns in service worker employment ρa_1 and the labor supply elasticity $wn_1'(w)/n_1(w)$ determine the productivity effect. The labor supply elasticity in the model aligns with the range of estimates reported in Lamadon et al. (2022). The IV estimate for the productivity effect is positive in the model at 0.02, though below the lower end of the 95% confidence interval in the data (0.04, 0.12). This difference highlights the tension between simultaneously matching the wage skill premium and the productivity effect. Our counterfactuals may thus underestimate the aggregate productivity effects of outsourcing.

5 Outsourcing, inequality and aggregate output

5.1 The rise in outsourcing

With our estimated model in hand, we are ready to ask how outsourcing affect aggregate output and inequality. We answer this quantitative question with two counterfactuals. The first counterfactual assesses how outsourcing affects the economy by contrasting the estimated 1997 economy to an economy without outsourcing similarly to Section 1. The economy without outsourcing equivalently corresponds either to an outright ban, or to an economy in which $\tau = 0$, $\eta^c = +\infty$ or $\overline{c} = +\infty$. This counterfactual does not require to take a stance on a particular driver of outsourcing beyond our estimated parameters, which we view as an advantage.

The second counterfactual evaluates how the rise in outsourcing reshapes the French labor market between 1997 and 2016. In that case, we have to take a stance on the underlying change in parameters. We choose a change in the entry cost of contractor η^c as our main specification because it is a neutral shock whose only direct effect is to increase how many contractors operate. We also explore changes in other parameters—job creation costs $\overline{c}$, comparative advantage τ, the correlation between productivity and outsourcing costs ι, and returns to scale ρ—in Table 6, Appendix C.

We compare steady-states of the estimated model. We select values of the entry cost η^c to match a particular aggregate outsourcing expenditure share, by which we index all our counterfactuals. As emphasized in Section 3.2, our firm-level expenditure data stops in 2008. To 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 2(b) (15%). In figures, we represent the economy without outsourcing and the baseline 1997 economy with circles, and subsequent

Figure 7: The effects of outsourcing on wages, earnings and welfare of service workers.

Note: 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 Online Appendix D.7. Circles: economy without outsourcing and baseline 1997 economy. Solid and dashed lines: 1997-2007 counterfactuals. Shaded lines: rise in outsourcing share from 0% to 6% if driven only by a change in η^c as between 1997 and 2007.

counterfactuals with lines.

5.2 Outsourcing and service workers

Our first quantitative result is that domestic outsourcing depresses the earnings and welfare of service workers. These effects crucially hinge on partially offsetting changes in employment and wages, which we unpack sequentially in Figure 7.

Domestic outsourcing has strong positive effects on service worker employment. Panel (a) indicates that their employment rate drops by 2.8 p.p. under an outsourcing ban relative to the 1997 baseline economy, and increases by 4.4 p.p. when outsourcing grows between 1997 and 2007. This increase in the employment rate of service workers coincides with a reallocation away from good producers and to contractors: the fraction of service workers employed at contractors rises by over 20 p.p. between 1997 and 2007, and by more than 40 p.p. through 2016.

The reallocation of low skill service workers toward contractors depresses wages because of the outsourcing wage penalty. Panel (b) shows that the average wage of service workers is 5% lower in 1997 than relative to an economy without outsourcing, and continues to fall as the outsourcing share reaches 11% in 2007. These wage losses are concentrated at the bottom of the wage distribution where contractors primarily operate. Outsourcing additionally removes the highest paying jobs from the labor market, leading to moderate wage declines in the top half of the distribution when outsourcing increases. The larger decline in wages at the bottom of the distribution implies that outsourcing increases within-skill inequality.

At the same time, our quantitative framework uncovers a nonlinear competition effect of outsourcing. When outsourcing grows between 1997 and 2007, wages in the bottom half of the distribution decline by close to 8%. As outsourcing continues to grow after 2007, competition for workers at the bottom of the job ladder—between contractors—intensifies. Wages in the bottom half of the wage distribution start to increase again. By 2016, wage losses diminish to 6%. This non-monotonicity highlights the interaction of reallocation and competition forces in general equilibrium and, by 2016, partially offsets the rise in within-skill inequality that opened through 2007.

How do these employment and wage effects shape earnings and welfare of low skill service workers? In our environment, expected earnings coincide with welfare, which we then exactly decompose into employment and wage contributions (see Online Appendix D.7 for details). Panel (c) shows that the reallocation of low skill service workers toward low-paying contractors depresses earnings by over 2.5% in the baseline economy with outsourcing relative to an economy without outsourcing. Similarly, rising outsourcing between 1997 and 2016 results in nearly 5% earnings losses. The simultaneous rise in low skill service employment is a strong general equilibrium feedback which largely offsets reallocation effects. It contributes 2% earnings losses under an outsourcing ban, and 4% earnings gains between 1997 and 2016. The non-monotonic general equilibrium response of the wage distribution thus closely approximates total expected earnings changes.

On net, low skill service workers lose from outsourcing. They experience 3.1% earnings and welfare gains when outsourcing is banned. When outsourcing rises between 1997 and 2007, their earnings and welfare decline by 1.7%. However, these losses shrink to 0.5% as the competition channels grows in importance by 2016. Thus, low skill service workers lose from outsourcing when outsourcing is prevalent, but not necessarily when it is omnipresent.

Our results indicate that general equilibrium effects are critical to evaluate how earnings and welfare of low skill service workers change when outsourcing grows in importance. The partial equilibrium impact alone would miss important general equilibrium margins of adjustment. Our model lets us quantify these margins and aggregate them into welfare.

5.3 Outsourcing and rent-sharing with service workers

We unpack the non-monotonic response of the wage distribution for service workers in Figure 8 given how central it is to our conclusions. Panel (a) shows average wages, wages at producers and wages at contractors. The outsourcing wage penalty is apparent as average producer wages are above contractor wages for the outsourcing shares we consider in our main analysis.[28] Panel (a) confirms that the non-monotonicity in wages in the bottom half of the distribution post 2007 is driven by contractors. By contrast, producer wages rise monotonically with outsourcing.

Did changes in rent-sharing cause the rise of producer wages and the non-monotonic changes in contractor wages? Panel (b) displays a simple decomposition of wages into markdowns and marginal products of labor according to the identity: $\mathbb{E}[\log w_s(z)] = \mathbb{E}[\log \mathrm{MPL}_s(z)] + \mathbb{E}[\log \mathrm{markdown}_s(z)]$, for service workers $s = 1$. With our convention, a higher markdown means less market power for employers.

We find that the average producer markdown remains largely constant. The rise in producer wages is driven by a rise in the marginal product of labor: as contractors hire more service workers, fewer

[28] For outsourcing shares between 0% and 4%, we obtain an average outsourcing wage premium because of a selection effect: only the most productive contractors operate, and thus they pay high wages. This effect is not present in our baseline model in Section 2 in which contractors are homogeneous. In addition, our results characterize the wage of the marginal goods producer relative to contractors conditional on outsourcing costs, not the unconditional average wage gap.

Figure 8: The effects of outsourcing on the wages of service workers.

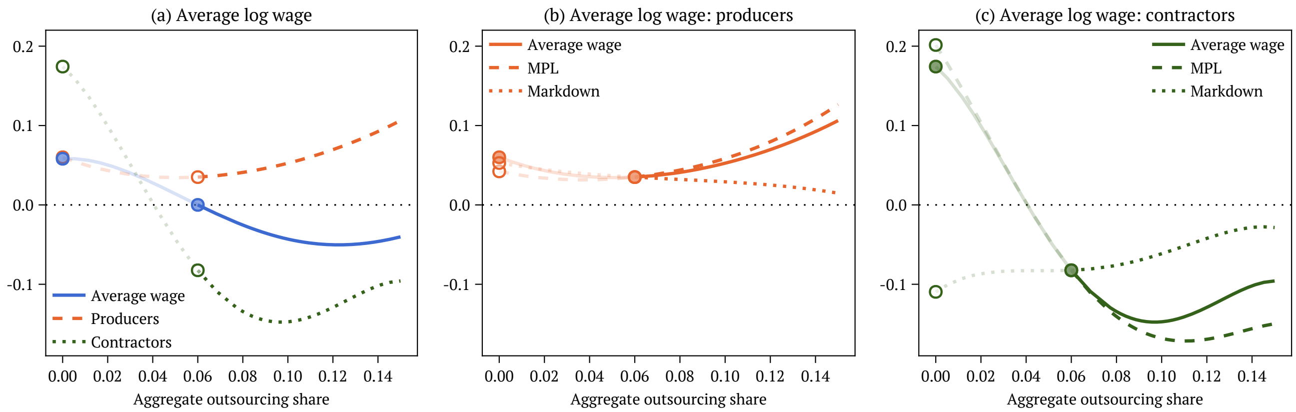

Note: All outcomes are shown as a function of aggregate outsourcing expenditures relative to the aggregate wage bill on the x-axis. Leftmost point corresponds to 1997 economy. Vertical dotted line corresponds to counterfactual 2007 economy. Rightmost point corresponds to counterfactual 2016 economy. MPL: Marginal Product of Labor. Circles: economy without outsourcing and baseline 1997 economy. Solid and dashed lines: 1997-2007 counterfactuals. Shaded lines: rise in outsourcing share from 0% to 6% if driven only by a change in η^c as between 1997 and 2007.

remain at goods producers, increasing the marginal product of labor. This effect crucially hinges on decreasing returns to scale and complementarities between worker types $\rho a_s < 1$.

In comparison, the non-monotonicity contractor wages reflects more nuanced forces. As outsourcing rises, markdowns steadily increase, reflecting a decrease in contractor market power due to the entry of more competitors. At the same time, the marginal product of labor at contractors falls through 2007 because of a general equilibrium force: the entry of contractors puts downward pressure on the price of outsourcing services, leading to lower revenue productivity at contractors. After 2007, a countervailing reallocation of workers toward the most productive contractors overcomes the decline in prices, leaving the marginal product of labor largely stable.

5.4 Outsourcing and non-service workers

We now turn to the effects of rising outsourcing on non-service workers to understand how between-skill inequality evolves. In contrast to service workers, outsourcing increases the expected earnings of low-skill core and high-skill workers through rising wages. Figure 9 displays our results. Panel (a) indicates that banning outsourcing has virtually no effect on non-service employment, although it mechanically reduces the employment share of goods producers that outsource. Rising outsourcing between 1997 and 2007 due to a decline in contractor entry costs increases non-service employment moderately, by just 0.1% by 2016. A substantial increase in the employment share of goods producers that outsource of more than 40 p.p. accompanies this modest increase in overall employment.

The strong response in the employment share at goods producers that outsource is key to understand the behavior of average wages in panel (b). In 1997, firms that outsource service workers pay nearly 5% higher wages to their non-service workers—a gap that largely persists throughout the period that we analyze. Two channels underlie this pay premium. The first channel is selection into outsourcing: goods producers that outsource are more productive on average and pay higher wages. The second channel is the productivity effect: outsourcing service workers enables goods producers to

Figure 9: The impact of outsourcing on non-service workers.

Note: 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 Online Appendix D.7. Circles: economy without outsourcing and baseline 1997 economy. Solid and dashed lines: 1997-2007 counterfactuals. Shaded lines: rise in outsourcing share from 0% to 6% if driven only by a change in η^c as between 1997 and 2007.

expand low-skill labor services, which raises the marginal product of labor, markdowns and wages for non-service workers through complementarities in production.

Together with the rise in non-service employment at firms that outsource, the pre-existing wage gap implies that average wages for non-service workers decline by nearly 2% under an outsourcing ban, and rise by nearly 2% between 1997 and 2016, as shown in panel (b). Panel (c) indicates that non-service workers lose over 2% earnings and welfare under an outsourcing ban, and gain over 2% between 1997 and 2007. Given moderate employment gains in panel (a), welfare and earnings are exclusively driven by wages. These gains for non-service workers contrast with losses experienced by service workers and contribute to rising inequality between skills.

5.5 Outsourcing, output and total factor productivity

Despite earnings losses for service workers, the productivity effect can lead to output gains in the aggregate. Indeed, we find that outsourcing substantially increases aggregate output. We describe the effects of outsourcing on aggregate output in Figure 10. Panel (a) indicates that banning outsourcing lowers output by 1.8%. Correspondingly, rising outsourcing after 1997 is associated with an output increase of 1% by 2007 and of 2% by 2016, in line with the productivity effects of outsourcing.

Do extensive or intensive margin gains underpin the rise in output? Panel (a) decomposes output changes into an extensive margin—employment—and an intensive margin—aggregate TFP. Aggregate employment effects drive the majority of the output gains between 1997 and 2007. These gains occurs almost exclusively for service workers. Under either the outsourcing ban or between 2007 and 2016, TFP gains contribute between one and two thirds of the changes.

Why does TFP remain relatively stable between 1997 and 2007? In principle, outsourcing allows productive firms with a high marginal product of labor to expand. Panel (b) reveals that the stability of TFP through 2007 masks two offsetting effects which we capture with an exact decomposition. We define two labor aggregators consistent with our production function. The first one is simply the

Figure 10: The impact of outsourcing on aggregate output and TFP.

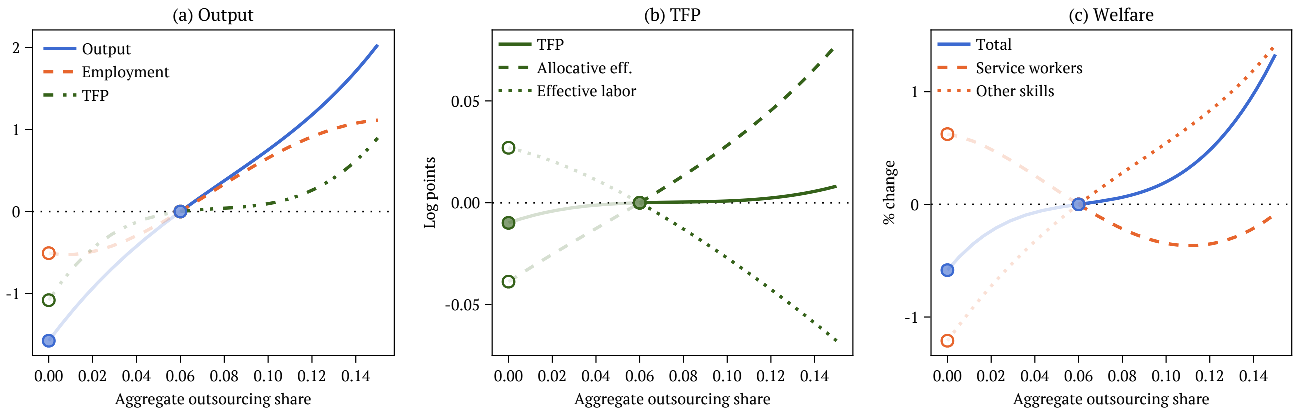

Note: All outcomes are shown as a function of aggregate outsourcing expenditures relative to the aggregate wage bill on the x-axis. Leftmost point corresponds to 1997 economy. Vertical dotted line corresponds to counterfactual 2007 economy. Rightmost point corresponds to counterfactual 2016 economy. Panel (b): TFP calculated as in equation (15). Circles: economy without outsourcing and baseline 1997 economy. Solid and dashed lines: 1997-2007 counterfactuals. Shaded lines: rise in outsourcing share from 0% to 6% if driven only by a change in η^c as between 1997 and 2007.

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

$$\Delta \log \text{TFP} = \underbrace{\Delta \log \frac{M^g \mathbb{E}[R(z, \boldsymbol{n})]}{\widetilde{N}}}_{\substack{\text{Allocative efficiency} \\ \text{given effective labor}}} + \underbrace{\Delta \log \frac{\widetilde{N}}{\widehat{N}}}_{\text{Effective labor}} . \tag{15}$$

Consistent with the reallocation of labor toward highly productive firms that were constrained by labor market frictions, allocative efficiency rises by 7% between 1997 and 2007. At the same time, the reallocation of service workers toward less productive contractors drags down TFP by nearly the same amount, leaving TFP essentially flat. Crucially, the comparative advantage view—if τ was counterfactually larger than one—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. Under an outsourcing ban and after 2007, allocative efficiency takes over the decline in effective labor, and TFP changes by 1%.

How are output gains distributed between firms and workers? Our setup features free entry, so that average ex-post profits exactly offset entry costs. As a result, the only welfare-relevant changes are earnings of workers. For this reason, the labor share remains largely constant as outsourcing changes (Table 6, Appendix C). Panel (c) shows that the aggregate output gains from outsourcing ultimately accrue to non-service workers. Service workers benefit while non-service workers lose from an outsourcing ban. Symmetrically, service workers lose while non-service workers gain from the change in outsourcing between 1997 and 2007. Service workers recoup part of their losses by 2016 due to the competition effect at the bottom of the job ladder.

Conclusion

This paper starts with a theory of domestic outsourcing. We argue 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 propose new reduced-form evidence that confirms the productivity and redistributive effects of outsourcing. Finally, equipped with a structurally estimated model, we show that outsourcing largely deteriorates labor market prospects for low skill service workers, while it increases aggregate output and between skill inequality.

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, Gancia, Gino, and Zilibotti, Fabrizio (July 2015). “Offshoring and Directed Technical Change”. *American Economic Journal: Macroeconomics* 7.3, pp. 84–122.

Antràs, Pol (Nov. 2003). “Firms, Contracts, and Trade Structure”. *The Quarterly Journal of Economics* 118.4, pp. 1375–1418.

Antràs, Pol, Fort, Teresa C., and Tintelnot, Felix (2017 2017). “The Margins of Global Sourcing: Theory and Evidence from U.S. Firms”. *American Economic Review* 107.9, pp. 2514–64.

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 (2025). “Technological Change and Domestic Outsourcing”. *Forthcoming at: Journal of Labor Economics.*

Berger, David, Herkenhoff, Kyle, and Mongey, Simon (Apr. 2022). "Labor Market Power". *American Economic Review* 112.4, pp. 1147–93.

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 (2023). "The Geography of Unemployment". *The Quarterly Journal of Economics* 138.3, pp. 1507–1576.

Bilal, Adrien, 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 (1999). "An Empirical Equilibrium Job Search Model with Search on the Job and Heterogeneous Workers and Firms". *International Economic Review* 40.4, pp. 1039–1074. (Visited on 04/18/2022).

— (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 (2022). "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, Yuhong, and Huang, Na (2019). "Stabilized Barzilai-Borwein Method". *Journal of Computational Mathematics* 37.6, pp. 916–936.

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 (May 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, Pascuel, and Schoefer, Benjamin (2023). "Paying Outsourced Labor: Direct Evidence from Linked Temp Agency-Worker-Client Data". *The Review of Economics and Statistics* 105.1, pp. 206–216.

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.

Engbom, Niklas and Moser, Christian (2022). “Earnings Inequality and the Minimum Wage: Evidence from Brazil”. *American Economic Review* 112.12, pp. 3803–47.

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. (2024). “The Reallocation Effects of Domestic Outsourcing”. *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 (Aug. 2020). “Bartik Instruments: What, When, Why, and How”. *American Economic Review* 110.8, pp. 2586–2624.

Gouin-Bonenfant, Émilien (2022). “Productivity Dispersion, Between-Firm Competition, and the Labor Share”. *Econometrica* 90.6, pp. 2755–2793.

Grossman, Gene and Helpman, Elhanan (2002). “Integration versus Outsourcing in Industry Equilibrium”. *The Quarterly Journal of Economics* 117.1, pp. 85–120.

— (2005). “Outsourcing in a Global Economy”. *The Review of Economic Studies* 72.1, pp. 135–159.

Grossman, Gene and Rossi-Hansberg, Esteban (Dec. 2008). “Trading Tasks: A Simple Theory of Offshoring”. *American Economic Review* 98.5, pp. 1978–97.

Grossman, Sanford and Hart, Oliver (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 (Apr. 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 (June 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. (May 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 (Jan. 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*.

Lhuillier, Hugo (2022). "The Local Root of Wage Inequality". *Working Paper*.

Melitz, Marc J. (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity". *Econometrica* 71.6, pp. 1695–1725.

Micco, Alejandro and Perez, Francisca (2024). "Costs and Benefits of Firms' Use of Outsourced Labor: A Structural Estimation". *SSRN Working Paper 4801865*.

Morchio, Iacopo and Moser, Christian (May 2024). "The Gender Pay Gap: Micro Sources and Macro Consequences". *NBER Working Paper 32408*.

Mortensen, Dale T. and Vishwanath, Tara (1991). "Information Sources and Equilibrium Wage Outcomes". *Northwestern University, Kellogg School of Management Discussion Paper 948*.

Muñoz, Mathilde (2023). "Trading Nontradables: The Implications of Europe's Job-Posting Policy". *The Quarterly Journal of Economics* 139.1, pp. 235–304.

Petrongolo, Barbara and Pissarides, Christopher A. (June 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. (June 1997). "The Growth of Temporary Services Work". *Journal of Economic Perspectives* 11.2, pp. 117–136.

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". In: *Identification and Inference for Econometric Models*. Ed. by Donald W.K. Andrews. New York: Cambridge University Press, pp. 80–108.

Topkis, Donald M. (1998). *Supermodularity and Complementarity*. Princeton University Press. ISBN: 9780691032443.

Appendix

A Proofs

A.1 Proof of Proposition 1

We consider continuous wage offer distributions $F_s(w)$. Then: $n_s(w) = \frac{(1+k_s)e_s}{[1+k_s(1-F_s(w))]^2}$, and $n_s'(w) = \frac{2kn_{0s}F_s'(w)}{[1+k_s(1-F_s(w))]^3}$. Profit-maximization in (2) becomes:

$$\pi(z) = \max_{v_s \in [0,1]^S, w_s} R(z, \{n_s(w)v_s\}_s) - \int w_s n_s(w_s) v_s ds, \tag{16}$$

where $v_s \in [0,1]$ captures that the size constraint $n_s \le n_s(w)$ may not bind. Start from the FOC for wages in (2). We obtain $R_{n_s} n_s' - n_s - w_s n_s' = 0$.[29] Differentiating the objective in (16) w.r.t. v_s and using the FOC for wages, we obtain $\partial \left(R(z, \{n_s(w)v_s\}_s) - \int w_s n_s(w_s) v_s ds\right) / \partial v_s = (R_{n_s} - w_s) n_s > 0$. Thus, firms are always at the corner $v_s = 1$. Hence, (2) coincides with: $\pi(z) = \max_{w_s} \Pi[z, \{w_s\}_s] \equiv R(z, \{n_s(w)\}_s) - \int w_s n_s(w_s) ds$. 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 wages, $F(w_s(z)) = \Gamma(z)$ and $n_s(z) \equiv n_s(w_s(z)) = \frac{(1+k_s)e_s}{M^g[1+k_s(1-\Gamma(z))]}$.

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: $\frac{d\left(n_s(w)w\right)}{dw}\Big|_{w=w_s(z)} = \frac{dR(z, \boldsymbol{n}_{-s}(w_s(z)), n_s(w))}{dw}\Big|_{w=w_s(z)} = \frac{\partial R}{\partial n_s}(z, \boldsymbol{n}(z)) \cdot n_s'(w_s(z))$, where $\boldsymbol{n}_{-s}$ denotes the vector $\boldsymbol{n}$ 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

$$n_s(z) w_s'(z) = n_s'(z)\big(R_{n_s}(z, \{n_t(z)\}_t) - w_s(z)\big). \tag{17}$$

Integrate over z subject to the boundary condition $w_s(\underline{z}) = \underline{w}_s$ to obtain the formula in Proposition 1.

A.2 Proof of Proposition 3

We start with the case without idiosyncratic outsourcing costs: $\varepsilon_s \equiv 1$ for all firms. First define the cost function: $C_s(n) = \min\left\{w_s(n)n, p_s n\right\} = n \min\left\{w_s(n), p_s\right\}$, where $w_s(n)$ is the inverse function of $n_s(w)$. Then rewrite firms' profit-maximization problem as: $\pi(z, \boldsymbol{\varepsilon}) = \max_{\{n_s\}_s} R\big(z, \{n_s\}_s\big) - \int C_s(n_s) ds$. As in Proposition 1, this profit function 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

[29] This equality implies $(R_{n_s} - w_s)n_s' = n_s > 0$. Thus, $R_{n_s} > w_s$ and $n_s' > 0$.

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 $z > \hat{z}_s$.

We then consider the case with idiosyncratic outsourcing costs. Write the problem of the firm as: $\pi(z, \boldsymbol{\varepsilon}) = \max_{n_s} R(z, \{n_s\}) - \int_s C_s(n_s, \varepsilon_s) ds$, where $C_s(n, \varepsilon) = n \min\{w_s(n), p_s \varepsilon\}$. 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.

Outsourcing expenditures on skill s are $\mathcal{E}_s(z, \boldsymbol{\varepsilon}) = p_s \varepsilon_s o_s(z, \boldsymbol{\varepsilon}) n_s(z, \boldsymbol{\varepsilon})$, and total outsourcing expenditures are $\mathcal{E}(z, \boldsymbol{\varepsilon}) = \int \mathcal{E}_s(z, \boldsymbol{\varepsilon}) ds = \int p_s \varepsilon_s o_s^*(z, \boldsymbol{\varepsilon}) n_s(z, \boldsymbol{\varepsilon}) ds$. Since o_s and n_s conditional on outsourcing are increasing in z, so are skill-s and total outsourcing expenditures.

A.3 Proof of Proposition 4

We first present our proof without idiosyncratic outsourcing costs. We then add outsourcing costs. We consider decisions around the outsourcing threshold $\hat{z}_s$. For notational simplicity, we omit dependence on $\hat{z}_s$ in this subsection since all functions are evaluated at this value. For instance, we denote by $n_{\text{s,in}}$ in-house size at $\hat{z}_s$, and by $n_{s,\text{out}}$ outsourced size at $\hat{z}_s$.

Proof that $\boldsymbol{n_{s,\text{in}}(\hat{z}_s) < n_{s,\text{out}}(\hat{z}_s)}$. Supermodularity implies that n_s is weakly increasing in z, regardless of whether workers are in-house or outsourced. Taking the left- and right-limits of n_s around the outsourcing threshold, we obtain that $n_{s,\text{out}} \geq n_{s,\text{in}}$ at the outsourcing threshold. The same argument implies that $n_{-s,\text{out}} \geq n_{-s,\text{in}}$.

Suppose for a contradiction that $n_{s,\text{out}} = n_{s,\text{in}}$. Then: $w_s(n_{s,\text{in}}) < \partial(n w_s(n))/\partial n|_{n=n_{s,\text{in}}}$ from $w_s(n)$ strictly increasing in n. Then from the F.O.C. for in-house goods producers: $\partial(n w_s(n))/\partial n|_{n=n_{s,\text{in}}} = R_{n_s}(n_{s,\text{in}}, n_{-s,\text{in}})$. Under our conjecture (for contradiction): $R_{n_s}(n_{s,\text{in}}, n_{-s,\text{in}}) = R_{n_s}(n_{s,\text{out}}, n_{-s,\text{in}})$. Using that size at other skills weakly increases: $R_{n_s}(n_{s,\text{in}}, n_{-s,\text{in}}) \leq R_{n_s}(n_{s,\text{out}}, n_{-s,\text{out}})$. From the F.O.C. for goods producers that outsource: $R_{n_s}(n_{s,\text{out}}, n_{-s,\text{out}}) = p_s$. From cost-minimization: $p_s \leq w_s(n_{s,\text{out}})$. This series of inequalities and equalities thus shows that: $w_s(n_{s,\text{in}}) < w_s(n_{s,\text{out}})$, and hence that $n_{s,\text{in}} < n_{s,\text{out}}$.

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

Proof that $\boldsymbol{R_{n_s}(\hat{z}_s, n_{s,\text{in}}(\hat{z}_s), n_{-s,\text{in}}(\hat{z}_s)) > R_{n_s}(\hat{z}_s, n_{s,\text{out}}(\hat{z}_s), n_{-s,\text{out}}(\hat{z}_s))}$. Given $\hat{z}_s$, consider the function $\phi(x) = \max_{n_{-s}} R(\hat{z}_s, x, n_{-s}) - p_s x - C_{-s}(n_{-s})$, where C_{-s} is the labor cost function for skills other than s. ϕ represents profits conditional on outsourced employment of skill s. By definition, ϕ is maximized at $n_{s,\text{out}}$. Since $n_{s,\text{in}} < n_{s,\text{out}}$ and ϕ inherits concavity from the joint problem, $\phi'(n_{s,\text{in}}) > 0$.

Using the envelope theorem, this inequality implies: $R_{n_s}(\hat{z}_s, n_{s,\text{in}}, n_{-s,\text{in}}) > p_s = R_{n_s}(\hat{z}_s, n_{s,\text{out}}, n_{-s,\text{out}})$ where the last equality follows from the F.O.C. at $n_{s,\text{out}}$.

Idiosyncratic outsourcing costs. The proof follows exactly the steps in without idiosyncratic costs once we condition on the vector $\{\varepsilon_s\}_s$.

A.4 Proof of Proposition 5

We consider the neutral or comparative disadvantage $\tau_s \leq 1$ case first. We then turn to the comparative advantage $\tau_s \gg 1$. We omit the dependence of equilibrium objects on $\boldsymbol{\varepsilon}$ when unambiguous.

A.4.1 Neutral or comparative disadvantage $\tau_s \leq 1$

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

$$\varepsilon_s p_s = R_{n_s}(\hat{z}_s, n_{s,\text{out}}(\hat{z}_s), \boldsymbol{n}_{-s,\text{out}}(\hat{z}_s)) \tag{18}$$

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

$$\begin{aligned}
&\left(p_s\tau_s - w - \frac{n_s(w)}{n_s'(w)}\right)\Big|_{w_s=w_{s,\text{in}}(\hat{z})} dw && (19)\\
\leq &\left(p_s - w - \frac{n_s(w)}{n_s'(w)}\right)\Big|_{w=w_{\text{in}}(\hat{z})} dw && (\text{use } \tau_s \leq 1)\\
= &\left(\frac{1}{\varepsilon_s} R_{n_s}(\hat{z}_s, n_{s,\text{out}}(\hat{z}), \boldsymbol{n}_{-s,\text{out}}(\hat{z})) - w - \frac{n_s(w)}{n_s'(w)}\right)\Big|_{w=w_{s,\text{in}}(\hat{z})} dw && (\text{use } (18))\\
< &\left(R_{n_s}(\hat{z}_s, n_{s,\text{in}}(\hat{z}), \boldsymbol{n}_{-s,\text{in}}(\hat{z})) - w - \frac{n_s(w)}{n_s'(w)}\right)\Big|_{w=w_{s,\text{in}}(\hat{z})} dw. && (\text{use Prop. } 4 \ \&\ \varepsilon_s \geq 1)
\end{aligned}$$

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 $w_s(\hat{z}_s) < p_s\varepsilon_s$. Directly follows from cost-minimization.

A.4.2 Comparative advantage $\tau_s \gg 1$

We consider the comparative advantage $\tau_s \gg 1$ case. It suffices to proving the stated results in the limit $\tau_s \uparrow +\infty$ to prove the threshold property. We first lay out the main arguments in our proof under the assumption of a single worker type, an isoelastic revenue function and no idiosyncratic outsourcing costs. We then extend this proof to a Cobb-Douglas revenue function: $R(z, \boldsymbol{n}) = \left(z \prod_{s=1}^S n_s^{a_s}\right)^\rho$, $\sum_{s=1}^S a_s = 1$; and with only skill $s = 1$ that can be outsourced and with idiosyncratic outsourcing costs. Finally, we relax the Cobb-Douglas assumption.

Single worker type and no idiosyncratic outsourcing costs. We start by assuming that the revenue function is isoelastic: $R(z,n) = zn^\rho$ with $0 < \rho < 1$. Market clearing for labor services writes: $M^g \int_{\hat{z}}^{\overline{z}} \left(\frac{\rho z}{p}\right)^{\frac{1}{1-\rho}} d\Gamma(z) = \tau M^c \int n_{\text{out}}(w) dF(w)$. Since in-house employment is always bounded below and above, $M^c \int n_{\text{out}}(w) dF(w) \equiv N_{\text{out}}$ remains bounded. Hence, to a leading order when $\tau \to \infty$, market clearing implies: $\int_{\hat{z}}^{\overline{z}} (\rho z)^{\frac{1}{1-\rho}} d\Gamma(z) = \frac{N_{\text{out}} M^c}{M^g} \tau p^{\frac{1}{1-\rho}} \sim \tau p^{\frac{1}{1-\rho}}$.

$\hat{z}$ 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: $c_1 \hat{z}^{\frac{1}{1-\rho}} p^{-\frac{\rho}{1-\rho}} \geq c_2 \hat{z} + c_3$ when τ is large enough, for bounded and non-vanishing functions of τ, $c_i(\tau)$. So if $p \to 0$, this identity implies $\hat{z} \to \underline{z}$, all firms outsource, and the indifference condition ceases to hold.

Now we are ready for a guess and verify. We guess that, as $\tau \to +\infty$, $p \to 0$. Then $\hat{z}$ hits $\underline{z}$. The integral in market clearing then becomes constant, and market clearing implies: $1 \sim \tau p^{\frac{1}{1-\rho}} \implies p \sim \tau^{\rho-1} \to 0$, and the guess is verified. This argument proves that $\tau p \sim \tau^\rho \to +\infty$: the marginal product of labor of contractors becomes infinite.

We now prove that the marginal goods producer pays wages below wages of contractors with a similar strategy as in the comparative disadvantage case. Suppose for a contradiction that $w_{\text{in}}(\hat{z})$ is in the support of the wage distribution of contractors. Then the wage F.O.C. holds for goods producer $\hat{z}$ at $w = w_{\text{in}}(\hat{z})$. The marginal benefit from a wage change dw for a contractor at $w = w_{\text{in}}(\hat{z})$ is, up to $n'(w)$, proportional to: $\left(p\tau - w - \frac{n(w)}{n'(w)}\right)\Big|_{w=w_{\text{in}}(\hat{z})} dw > (p\tau - R_n(\hat{z}, n_{\text{in}}(\hat{z})))\, dw$. The marginal product of labor is bounded above by a constant independent from equilibrium objects: $R_n(\hat{z}, n_{\text{in}}(\hat{z})) \leq \overline{z} n_{\text{in}}(\underline{z})^\rho$. As $\tau \to +\infty$, $p\tau \sim \tau^\rho \to +\infty$. Combining both observations implies: $\left(p\tau - w - \frac{n(w)}{n'(w)}\right)\Big|_{w=w_{\text{in}}(\hat{z})} dw > 0$ as $\tau \to +\infty$.

Cobb-Douglas revenue function, only skill $s = 1$ can be outsourced, and idiosyncratic outsourcing costs. We now generalize the arguments above to a revenue function that is Cobb-Douglas, when only skill $s = 1$ can be outsourced and in the presence of idiosyncratic outsourcing costs. Market clearing for labor services writes:

$$M^g \left(\frac{\rho a_1}{p}\right)^{\frac{1}{1-\rho a_1}} \int_{\underline{\varepsilon}_1}^{\overline{\varepsilon}_1} \int_{\hat{z}(\varepsilon_1)}^{\overline{z}} \varepsilon_1 (z/\varepsilon_1)^{\frac{1}{1-\rho a_1}} \tilde{R}(z,\varepsilon_1)^{\frac{1}{1-\rho a_1}} d\Gamma(z,\varepsilon_1) = \tau \times M^c \int n_{\text{out}}(w) dF(w) \equiv \tau N_{\text{out}},$$

where we define $\tilde{R}(z,\varepsilon_1)$ for goods producers who outsource skill $s = 1$ such that: $R(z,\varepsilon_1) = z(n_1^{\text{out}})^{\rho a_1} \tilde{R}(z,\varepsilon_1)$, $\tilde{R}(z,\varepsilon_1) \equiv \prod_{s>1} (n_s^{\text{out}}(z,\varepsilon_1))^{\rho a_s}$. As above, N_{out} is bounded below and above. Thus, to a leading order: $\int_{\underline{\varepsilon}_1}^{\overline{\varepsilon}_1} \int_{\hat{z}(\varepsilon_1)}^{\overline{z}} \varepsilon (z/\varepsilon_1)^{\frac{1}{1-\rho a_1}} \tilde{R}(z,\varepsilon_1)^{\frac{1}{1-\rho a_1}} d\Gamma(z,\varepsilon_1) \sim \tau p^{\frac{1}{1-\rho a_1}}$.

$\tilde{R}(z,\varepsilon)$ only depends on in-house employment and so is bounded below and above using the minimum and the maximum of the in-house labor supply curve, which are independent from τ, p as they only depend on matching rates: $\underline{\mathcal{R}} \equiv \prod_{s>1} \underline{n}_s^{\rho a_s} \leq \tilde{R}(z,\varepsilon) \leq \prod_{s>1} \overline{n}_s^{\rho a_s} \equiv \overline{\mathcal{R}}$. Thus, to a leading order: $\int_{\underline{\varepsilon}_1}^{\overline{\varepsilon}_1} \int_{\hat{z}(\varepsilon_1)}^{\overline{z}} \varepsilon (z/\varepsilon_1)^{\frac{1}{1-\rho a_1}} d\Gamma(z,\varepsilon_1) \sim \tau p^{\frac{1}{1-\rho a_1}}$.

We are left with characterizing how $\hat{z}(\varepsilon_1)$ depends on τ and p. When the indifference threshold is interior to the support of the productivity distribution, $\hat{z}(\varepsilon_1)$ is defined by the indifference condition $R - nR_n + n^2 w'(n)\big|_{\text{in}} \leq R - nR_n\big|_{\text{out}}$. Specifically: $(1-\rho a_1) R\big|_{\text{in}} + n_{1,\text{in}}^2 w_1'(1,\text{in}) \leq (1-\rho a_1) R\big|_{\text{out}}$. Using the structure of the revenue function, we obtain:

$$(1-\rho a_1)\hat{z}n_{1,\text{in}}^{\rho a_1}\tilde{R}_{\text{in}} + n_{1,\text{in}}^2 w_{1,in}' \le (1-\rho a_1)\left(\frac{\rho a_1}{p}\right)^{\frac{\rho a_1}{1-\rho a_1}} \left(\hat{z}\varepsilon_1^{-\rho a_1}\right)^{\frac{1}{1-\rho a_1}} \tilde{R}_{\text{out}}^{\frac{1}{1-\rho a_1}}. \tag{20}$$

where we omit dependence on $\hat{z}(\varepsilon)$ and ε for notational simplicity. Thus, as $\tau \to \infty$, we obtain an equation of the form: $A\hat{z} + B \le C\left(\hat{z}/(p\varepsilon_1)^{\rho a_1}\right)^{\frac{1}{1-\rho a_1}}$, where A, B, C have finite limits.

We now guess and verify that as $\tau \to +\infty$, $p \to 0$. In that limit, outsourcing skill 1 is always preferable, and $\hat{z}(\varepsilon) \to \underline{z}$ until it hits $\underline{z}$. At that point, the indifference condition ceases to hold. The integral in the market clearing condition becomes a constant, and hence $\tau p^{\frac{1}{1-\rho a_1}} \sim 1$. Hence $p \sim \tau^{-(1-\rho a_1)} \to 0$ and the guess is verified.

A similar wage deviation argument to the one in the single worker type case ensures that marginal goods producers post wages below contractors.

General revenue function, only skill $s = 1$ can be outsourced, and idiosyncratic outsourcing costs. We now extend the proof to the a general revenue function under our strict Inada condition. For skill 1: $\frac{n_1 R_{n_1}}{R} \le \rho < 1$. Integrating this condition (seeing it as $d\log R/d\log n_1$) we immediately get that there is a function $\mathcal{R}(z, \varepsilon_1)$—that depend on the optimal in-house hiring in other skills, that remains bounded—such that: $R(z, n) \le \mathcal{R}(z, \varepsilon_1)n_1^{\rho}$.

We consider labor services market clearing. Demand for outsourced labor satisfies: $R_{n_1}(z, \varepsilon_1) = p_1\varepsilon_1$. Using the inequality above: $p_1\varepsilon_1 n_1 = \frac{n_1 R_{n_1}}{R} \times R \le \rho\mathcal{R}(z, \varepsilon_1)n_1^{\rho}$. Hence, we bound: $n_1 \le \left[\frac{\rho\mathcal{R}(z,\varepsilon_1)}{p_1\varepsilon_1}\right]^{\frac{1}{1-\rho}}$. We then extend the derivation using the market clearing condition exactly as in the Cobb-Douglas case: $\tau N_{\text{out}} \le M^g \left(\frac{\rho}{p_1}\right)^{\frac{1}{1-\rho}} \int_{\underline{\varepsilon}_1}^{\overline{\varepsilon}_1} \int_{\hat{z}(\varepsilon_1)}^{\overline{z}} \varepsilon_1^{-\frac{\rho}{1-\rho}} \mathcal{R}(z, \varepsilon_1)^{\frac{1}{1-\rho}} d\Gamma(z, \varepsilon_1)$. Thus, we have shown the existence of a bounded functional $I(\hat{z})$ such that: $\tau p^{\frac{1}{1-\rho}} \le I(\hat{z})$.

Next, using that $(1-\rho)R \le R - nR_n \le R$, we obtain (20) as under Cobb-Douglas revenues, and thus: $\hat{z} \to \underline{z}$ as $p \to 0$. Hence, as under under Cobb-Douglas revenues we guess that $p \to 0$ as $\tau \to \infty$. We verify the guess using our bound with I. Wages then follow as under Cobb-Douglas revenues.

Proof that $\boldsymbol{w_1(\hat{z}_1) < p_1\varepsilon_1}$. Again follows from cost-minimization.

A.5 Proof of Corollary 1

The hire rate from employment is increasing in w, while churn is decreasing in w. The comparative statics in Proposition 1 for these outcomes immediately follow from Proposition 5.

To show that net poaching is increasing in w as well, combine equations (21) and (22) 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 D.8, $q_s = \frac{e_s\delta_s(1+k_s)}{\widehat{M}_s}$. Hence $\text{NP}_s(w) = (1-\phi_s)F_s(w)(\delta_s+\lambda_s^E(1-F_s(w))) - \lambda_s^E(1-F_s(w)) = (1-\phi_s)\delta_s F_s(w) - \lambda_s^E(1-F_s(w))\big[1-(1-\phi_s)F_s(w)\big]$. Viewed as a function of $F_s(w)$, this quantity is weakly increasing in $F_s(w)$ if the second component $F \mapsto (1-F)\big[1-(1-\phi)F\big]$ 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 5.

B Reduced-form results

Table 3: Dependent variable: log spending on external workers.

	All			Exporters				
	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS
Log V.A. + Out.	1.07***	1.11***	0.83***	0.81***	0.99***			
	(0.01)	(0.01)	(0.02)	(0.02)	(0.21)			
Log V.A.						0.98***		
						(0.22)		
Log Size							1.66***	
							(0.34)	
Log Labor Prod.								2.41***
								(0.67)
Fixed Effects								
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 4: Dependent variable: spending on external workers as a fraction of labor costs, in p.p.

	All			Exporters				
	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS
Log V.A. + Out.	1.73***	1.82***	2.63***	1.85***	3.34**			
	(0.03)	(0.03)	(0.08)	(0.12)	(1.07)			
Log V.A.						3.41**		
						(1.09)		
Log Size							6.04***	
							(1.82)	
Log Labor Prod.								8.48**
								(2.98)
Fixed Effects								
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.

Table 5: The productivity effect of outsourcing.

	First stage	Reduced form		2SLS	
	Δ out. share	Δ log R	Δ log VA	Δ log R	Δ log VA
Change in out. IV	0.207***	0.019***	0.018***		
	(0.044)	(0.003)	(0.003)		
Change in out. share				0.092***	0.086***
				(0.021)	(0.022)
Obs.	47823	47823	47823	47823	47823
1^{st}-stage F-stat.					22.150

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 4.

C Unpacking the rise in outsourcing

Section 5 focuses on changes in the contractor entry cost η^c as the driver of the rise in outsourcing post 1997. Of course, other changes in the economy may also have contributed to this rise. We explore the role of changes in other parameters in Table 6. We consider changes in job creation costs $\overline{c}$, comparative advantage τ, the correlation between productivity and outsourcing costs ι, and returns to scale ρ. We construct counterfactuals that change one parameter at a time such that the increase in the aggregate outsourcing share is 5 p.p. as in our baseline analysis.

Changes in each of these parameters have broadly similar effects across multiple outcomes. In each case, the overall employment rate and that of service workers increases. The employment share of contractors among all workers and among service workers also rise. Value added increases for all parameters except decreasing returns, as the model requires a decline in ρ to generate rising outsourcing. The labor share remains relatively stable for most parameters.

There are more meaningful differences for wages and earnings. Changes in entry costs η^c, job creation costs $\overline{c}$ and returns to scale ρ all imply substantial declines in service wages, and therefore earnings losses. An increase in comparative advantage τ implies a much more muted service wage response because the marginal product of labor at contractors rises at the same time as service workers reallocate there. Hence, service wages and earnings ultimately rise. An increase in demand through a decline in the correlation ι also implies wage and earnings gains for service workers, albeit through the equilibrium price of outsourcing rather than comparative advantage. Both channels operate through the marginal revenue product at contractors and the upward-sloping labor supply curve.

Table 6: The aggregate consequences of outsourcing under various shocks.

	Ban	Entry (η^C)	Vacancy ($\bar{c}^C$)	Prod. (τ)	Demand (ι)	DRS (ρ)
Employment rate (p.p.)						
Total	-0.67	0.92	0.93	0.30	0.38	0.06
Low-skill service	-2.81	4.35	4.40	1.23	1.40	1.00
Other skills	-0.12	0.03	0.03	0.06	0.12	-0.19
Contractor employment share (p.p.)						
Total	-5.59	6.16	6.17	4.99	4.76	4.88
Low-skill service	-27.3	28.0	28.0	23.6	22.5	23.1
Average wage (%)						
Total	-0.24	-0.24	-0.27	1.50	1.75	-4.67
Low skill service	5.14	-4.54	-4.65	-0.52	0.59	-6.21
Other skills	-1.68	0.92	0.91	2.03	2.06	-4.26
Value added (%)	-1.78	1.00	0.98	2.17	2.60	-3.99
TFP (%)	-1.15	0.16	0.14	1.90	2.26	-4.72
Labor share (p.p.)	0.53	-0.15	-0.15	-0.24	-0.29	-0.45
Worker expected earnings (%)						
Total	-0.67	0.36	0.33	1.62	1.91	-4.46
Low-skill service	3.09	-1.72	-1.80	0.28	1.46	-5.38
Other skills	-1.69	0.90	0.89	1.98	2.05	-4.20

Note: Counterfactuals constructed by changing one parameter at a time η^c, $\bar{c}$, τ, ι or ρ to match a 5 p.p. increase in the aggregate outsourcing share from baseline 1997 economy with a 6% aggregate outsourcing share.

Online Appendix

Outsourcing, Inequality and Aggregate Output

Adrien Bilal Hugo Lhuillier

D Additional proofs

D.1 Workers and the labor supply curve

This section 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 simply maximize income: they always accept higher wage offers while employed. Equating the value of being employed to the value of being unemployed defines the reservation wage $\underline{w}_s$, given in Online Appendix D.2.

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 $G_s(w)$:

$$G_s(w) = \frac{F_s(w)}{1 + k_s(1 - F_s(w))}, \qquad k_s = \frac{\lambda_s^E}{\delta_s}. \tag{21}$$

From equation (21) we characterize the number $N_s(w)$ of employed workers per wage offer w for every skill s:

$$N_s(w) = \frac{(1 + k_s)e_s}{\big(1 + k_s(1 - F_s(w))\big)\big(1 + k_s(1 - F_s(w^-))\big)}, \tag{22}$$

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.

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: $\lambda_s^U F_s(w) u_s = \big(\delta_s + \lambda_s^E(1 - F_s(w))\big)(m_s - u_s)G_s(w)$, 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 (21).

$N_s(w)$ is equal to the limit of the ratio $\frac{G_s(w) - G_s(w-\varepsilon)}{F_s(w) - F_s(w-\varepsilon)}$ when $\varepsilon \to 0$, times the number of employed workers $m_s - u_s$. Straightforward differentiation delivers (22).

D.2 Reservation wage

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

$$\lambda^U \underline{w} = \lambda^E b + (\lambda^U - \lambda^E) \left[b + \lambda^U \int_{\underline{w}}^\infty \frac{(1 - F(x))dx}{r + \delta + \lambda^E(1 - F(x))} \right]. \tag{23}$$

D.3 CES revenue function

Consider the revenue function $R(z, \boldsymbol{n}) = z \left(\int (a_s n_s)^{\frac{\eta-1}{\eta}} ds \right)^{\frac{\eta}{\eta-1} \frac{\sigma-1}{\sigma}}$. 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.

To verify supermodularity, we calculate:

$$\begin{aligned} \frac{\partial^2 R}{\partial z \partial n_k} &= \frac{\sigma - 1}{\sigma} \left(\int (a_s n_s)^{1 - \frac{1}{\eta}} ds \right)^{\frac{\eta(\sigma-1)}{\sigma(\eta-1)} - 1} a_k^{1 - \frac{1}{\eta}} n_k^{-\frac{1}{\eta}} > 0 \\ \left. \frac{\partial^2 R}{\partial n_k \partial n_\ell} \right|_{k \neq \ell} &= \frac{\sigma - 1}{\sigma} \left(\frac{\sigma - 1}{\sigma} - \frac{\eta - 1}{\eta} \right) \left(\int (a_s n_s)^{1 - \frac{1}{\eta}} ds \right)^{\frac{\eta(\sigma-1)}{\sigma(\eta-1)} - 2} a_k^{1 - \frac{1}{\eta}} a_\ell^{1 - \frac{1}{\eta}} n_k^{-\frac{1}{\eta}} n_\ell^{-\frac{1}{\eta}}. \end{aligned}$$

The second line is positive if $\sigma > \eta$.

By comparing the curvature in the revenue function 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.

This CES revenue function satisfies our strict Inada condition:

$$\frac{n_s R_{n_k} dk}{R} = \frac{\sigma - 1}{\sigma} \frac{n_k^{\frac{\eta-1}{\eta}} dk}{\int n_s^{\frac{\eta-1}{\eta}} ds} < \frac{\sigma-1}{\sigma} < 1.$$

D.4 Proof of Proposition 2

Existence and uniqueness among equilibria with continuous F. Proposition 1 suffices to complete a guess and verify strategy to exhibit an equilibrium with a continuous wage offer distribution. There are two conditions to verify.

1. Reservation wage. The first condition to verify is whether a reservation wage compatible with those results exists. Omitting s subscripts, re-write (23) as

$$\lambda^U \underline{w} = \lambda^E b + (\lambda^U - \lambda^E) \left[b + \lambda^U \int_{\underline{z}}^{\infty} \frac{(1 - \Gamma(x)) w'(x) dx}{r + \delta + \lambda^E (1 - \Gamma(x))} \right]. \tag{24}$$

$w'(x)$ is a function of the reservation wage $\underline{w}$ through the ODE (17). 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 (17). Differentiating (17), 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 (24) is a decreasing function of $\underline{w}$ that goes to $\lambda^E b \le \lambda^U b$ 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 $\underline{w}$.

2. No mass points. The second condition to verify is that in equilibrium $R_{n_s} > w_s$, so that we can take the FOC w.r.t. wages, that firms are always at the corner $v_s = 1$ and that there are no mass points. Integrating by parts the wage expression in Proposition 1, we obtain:

$$n_s(z)\big(R_{n_s}(z, \boldsymbol{n}(z)) - w_s(z)\big) = n_s(\underline{z})\big(R_{n_s}(\underline{z}, \boldsymbol{n}(\underline{z})) - \underline{w}\big) + \int_{\underline{z}}^{z} \frac{dR_{n_s}(x, \boldsymbol{n}(x))}{dx} n_s(x) dx. \tag{25}$$

We may rewrite: $\frac{dR_{n_s}(z, \boldsymbol{n}(z))}{dz} = \Xi\left(\xi + (\theta - 1)\phi\right)$, where $\Xi = \frac{R_{n_s}}{z}$ is the ratio of the marginal product of labor to productivity, $\xi \equiv \frac{z R_{n_s z}}{R_{n_s}}$ is the elasticity of the marginal product of labor w.r.t. productivity, $\theta - 1 \equiv \frac{n_s R_{n_s n_s}}{R_{n_s}}$ is the elasticity of the marginal product of labor w.r.t. employment, and $\phi \equiv \frac{z n_s'(z)}{n_s(z)} = \frac{2k_s z \Gamma'(z)}{1 + k_s(1 - \Gamma(z))}$ is the elasticity of employment w.r.t. productivity. Without further restrictions, all these elasticities depend on z and $\boldsymbol{n}(z)$.

To gain intuition, consider the Cobb-Douglas case we use in Sections 4 and 5. In that case, up to renormalizing productivity $z^\rho \equiv z$, $\Xi = \rho a_s n_s(z)^{\rho a_s - 1} \prod_{k \neq s} n_k(z)^{\rho a_k}$. Since $n_k(z), n_s(z)$ are bounded above and below by constants, Ξ is also bounded above and below by constants. Similarly, in that case, $\xi = 1$, $\theta = \rho a_s$, $\phi = \frac{2k_s z \Gamma'(z)}{1 + k_s(1 - \Gamma(z))}$ is bounded between 0 and a positive constant since it is a continuous function on an compact interval. We denote by $n_s^0 = \frac{(1+k_s)e_s}{M^g}$ and generalize these insights by imposing the following assumption.

Assumption (A'). There exists a constant $X > 1$ independent from $\underline{z}, \overline{z}$ such that $|\Xi(z,\boldsymbol{n})|$, $|\xi(z,\boldsymbol{n})|$, $|\theta(z,\boldsymbol{n})|$, $|\phi(z)| \leq X$ and $|\Xi(z,\boldsymbol{n})| \geq 1/X$ for all $\boldsymbol{n}$ such that $n_s \in \left[n_s^0/(1+k_s), n_s^0\right]$, and for all $z \in [\underline{z}, \overline{z}]$.

Assumption (A') is a technical regularity condition to ensures that the revenue function is well-behaved. It rules out extreme cases in which productivity or employment changes have particularly strong effects on the marginal product of labor or the shape of the productivity distribution. The revenue function we use for our quantification satisfies this condition.

Under Assumption (A'), $\left|\int_{\underline{z}}^{z} \frac{dR_{n_s}(x,\boldsymbol{n}(x))}{dx} n_s(x)dx\right| \leq I|z-\underline{z}|$ for a constant $I > 0$ independent from $\underline{z}, \overline{z}$. We now consider how $R_{n_s}(z,\boldsymbol{n}(z)) - w_s(z)$ changes as $\underline{z} \uparrow +\infty$ and $|\lambda^E - \lambda^U| \downarrow 0$. The reservation wage expression (24) rewrites, after substituting in the solution to the wage ODE:

$$\begin{aligned}
\underline{w}_s &= \frac{\lambda_s^E}{\lambda_s^U} b_s + (\lambda_s^U - \lambda_s^E)\left[\frac{b_s}{\lambda_s^U} + \frac{1}{\delta_s}\int_{\underline{z}}^{\overline{z}} \frac{(1-\Gamma(x))}{1+k_s(1-\Gamma(x))}\Big(R_{n_s}(x,\boldsymbol{n}(x)) - w_s(x)\Big)\frac{n_s'(x)}{n_s(x)}\right] \\
&\equiv c_1 + c_2(\lambda_s^U - \lambda_s^E)\left[\int_{\underline{z}}^{\overline{z}} \frac{(1-\Gamma(x))\phi(x)}{1+k_s(1-\Gamma(x))}\left(\frac{R_{n_s}(x,\boldsymbol{n}(x)) - w_s(x)}{x}\right)dx\right],
\end{aligned}$$

for constants $c_1, c_2 > 0$. (25) then implies:

$$n_s(z)\frac{R_{n_s}(z,\boldsymbol{n}(z)) - w_s(z)}{z} = n_s(\underline{z})\frac{R_{n_s}(\underline{z},\boldsymbol{n}(\underline{z})) - \underline{w}_s}{\underline{z}}\frac{\underline{z}}{z} + \frac{1}{z}\int_{\underline{z}}^{z} \frac{dR_{n_s}(x,\boldsymbol{n}(x))}{dx} n_s(x)dx,$$

so that $\left|\frac{R_{n_s}(z,\boldsymbol{n}(z)) - w_s(z)}{z}\right| \leq c_3\left|\frac{R_{n_s}(\underline{z},\boldsymbol{n}(\underline{z})) - \underline{w}_s}{\underline{z}}\right| + c_4$ for constants $c_3, c_4 > 0$. Substituting back into the reservation wage equation, and since $\frac{(1-\Gamma(x))\phi(x)}{1+k_s(1-\Gamma(x))}$ is bounded above, we obtain that $\underline{w}_s$ remains bounded as $\underline{z} \uparrow +\infty$.

Hence, $R_{n_s}(\underline{z},\boldsymbol{n}(\underline{z})) - \underline{w}_s \uparrow +\infty$ as $\underline{z} \uparrow +\infty$. Then (25) together with Assumption (A') implies that $R_{n_s}(z,\boldsymbol{n}(z)) - w_s(z) \uparrow +\infty$ as $\underline{z} \uparrow +\infty$. Thus, we have shown that for high enough $\underline{z}$, the marginal product of labor is always above the wage, and our guess of an equilibrium with a smooth wage offer distribution is verified. Uniqueness among smooth equilibria immediately follows from Proposition 1 and the unique solution for the reservation wage.

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 (2) becomes:

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

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 $\mathcal{C}^\infty$ 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: $|H_\sigma(\varepsilon)-\mathbb{1}\{\varepsilon \geq 0\}| \leq h_0\sigma$, 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 $\sigma \to 0$. Fourth, we show that wages converge to our candidate equilibrium when $\sigma \to 0$.

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 $\mathcal{C}^\infty$ 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.1. 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 $F(w(z_0)) = \mathbb{P}[\varepsilon \leq w(z_0) - w(z)] = \int H_\sigma\big(w(z_0)-w(z)\big)d\Gamma(z)$. Then

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

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

4. Wages when $\sigma \downarrow 0$. We go back to the maximization problem (16) and use an argument that resembles Berge's maximum theorem that we cannot apply directly. Re-write (16) as choosing the wage $w_\sigma(Z)$ of a firm with productivity Z. The wage function w_σ must hence satisfy $z = \operatorname{argmax}_Z R(z, n(w_\sigma(Z))) - w_\sigma(Z)n(w_\sigma(Z))\big]$. 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 $n(w_\sigma(Z)) \to n(Z)$, 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

$z = \text{argmax}_Z\, R(z, n(Z)) - w_0(Z)n(Z)$. 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.

D.5 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

$$\pi^C(w) = \max_k p_s k^{1-\beta} n_s(w)^\beta - r_s k - w n_s(w). \tag{26}$$

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 (26), 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 $\beta \to 1$. Then, (26) becomes $\pi^C(w) = (\tau_s p_s - w)n_s(w)$.

D.6 Outsourcing equilibrium

We focus on the case without idiosyncratic outsourcing costs $\varepsilon_s \equiv 1$. 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 $z \in [\underline{z}, z_{1s}]$ 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 equilbrium size are

$$\Upsilon_{1s}(z) = \frac{M^g\Gamma(z)}{\widehat{M}_s}, \quad n_{1s}(z) = \frac{n_{0s}}{\widehat{M}_s[1 + k_s(1 - \Upsilon_{1s}(z))]^2}, \quad z \in [\underline{z}, z_{1s}]. \tag{27}$$

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 1 stems from the number of workers they attract and retain with a wage offer given in equations (27). 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}: $R_{n_s}(z_{1s}, \boldsymbol{n}(z_{1s})) = p_s$.

In the intermediate wage region, contractors compete with in-house goods producers. Contractors being homogeneous, they 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:

$$F_{2s}(w) = 1 + \frac{1}{k_s}\left[1 - \left(\frac{n_{0s}}{\widehat{M}_s n_{2s}(w)}\right)^{\frac{1}{2}}\right], \quad n_{2s}(w) = \frac{p_s - w(z_{1s})}{p - w} n_{1s}(z_{1s}), \quad z \in [z_{1s}, z_{2s}]. \tag{28}$$

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

$$\Upsilon_{3s}(z) = \frac{M^c + M^g\Gamma(z)}{\widehat{M}_s}, \quad n_{3s}(z) = \frac{n_{0s}}{\widehat{M}_s[1 + k_s(1 - \Upsilon_{3s}(z))]^2}, \quad z \in [z_{2s}, \hat{z}_s]. \tag{29}$$

Starting from $w_s(z_{2s})$, wages once again follow Proposition 1 but with an equilibrium size given in equations (29). Above the threshold productivity $\hat{z}_s$, no firm operates in-house. There, firms outsource their employment.

D.7 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 $r \to 0$. We show that, when $r \to 0$, 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 $V(x) = \int xr \int_0^\infty e^{-rt} h(t, x) d\mu(x)$ 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 $r \int_0^\infty e^{-rt} u(t) dt \to \lim_{t\to\infty} u(t)$ for any smooth and bounded function u, we can apply this last result x by x and obtain $V(x) \to \int xh(\infty, x)dx \equiv \mathbb{E}[x]$. To show $r \int_0^\infty e^{-rt} u(t) dt \to \lim_{t\to\infty} u(t)$ for any smooth function u, change variables $\tau = rt$: $r \int_0^\infty e^{-rt} u(t) dt = \int_0^\infty e^{-\tau} u(\tau/r) d\tau$. $u(\tau/r) \to u(\infty)$ 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 $\overline{w} = \frac{e^G w^G + e^C w^C}{e^G + e^C}$ the average wage in the economy, and $\overline{e} = e^G + e^C = 1 - u$ 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

$$\Delta I = \underbrace{(\overline{w} - b)\Delta\overline{e}}_{\text{Employment}} + \underbrace{\left(\frac{(e^G)'}{\overline{e}'}\Delta w^G + \frac{(e^C)'}{\overline{e}'}\Delta w^C\right)}_{\text{Wage}} + \underbrace{(w^C - w^G)\Delta\left(\frac{e^C}{\overline{e}}\right)}_{\text{Wage penalty}}. \tag{30}$$

D.8 Dynamic firm problem

We first show that the size constraint in (2) 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. We consider the case in which a single firm faces a dynamic problem, while all other firms are at their stationary behavior. 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 (21), $\phi + (1-\phi)G(w) = \frac{1}{1+k(1-F(w))}$, 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 (2) 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 can pay different wages to new workers, but face an equal-pay constraint within worker type and within time periods. 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 $r \to 0$,

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

Substituting into the first-order condition, we obtain

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

which coincides with the static first-order condition.

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 substracting 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 (2023) 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.

Summary statistics. Table 7 presents summary statistics of our firm-level dataset (FICUS-EAE-DADS *Postes*-Trade). Table 8 presents summary statistics of our worker-level dataset (DADS *Panel*).

Table 7: Firm-level dataset summary statistics.

	Median	Mean	10th perc.	90th perc.	Count
Value added (1,000 euros)	1,311	5,851	224	7,612	216,051
Contractor	1,442	8,608	179	10,356	8,207
Goods producer	1,307	5,742	226	7,492	207,844
Employment	32	97	6	148	216,051
Contractor	43	279	5	311	8,207
Goods producer	31	89	6	141	207,844
Payroll (1,000 euros)	708	2,605	125	3,680	216,051
Contractor	916	5,640	112	6,525	8,207
Goods producer	702	2,485	125	3,578	207,844
Outsourcing exp. (1,000 euros)	9	231	0	384	216,051
Contractor	0	168	0	247	8,207
Goods producer	10	234	0	389	207,844
Exporter status	0	0	0	1	216,051
Contractor	0	0	0	0	8,207
Goods producer	0	0	0	1	207,844

Table 8: Worker-level dataset summary statistics.

	Median	Mean	10th perc.	90th perc.	Count
Employment rate	1.00	0.85	0.00	1.00	13,167,711
Non-service worker	1.00	0.85	0.00	1.00	12,803,934
Service worker	1.00	0.76	0.00	1.00	363,777
Annualized wage (1,000 euros)	27.53	33.80	16.10	56.51	13,167,711
Non-service worker	27.72	34.10	16.29	57.05	12,803,934
Service worker	21.63	23.29	10.08	36.72	363,777

F Reduced-form results

F.1 Selection into outsourcing

Table 9: Summary statistics

Panel A: Outsourcers vs. non-outsourcers				
	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

Panel B: Outsourcing by industry		
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 11: Outsourcing by value added.

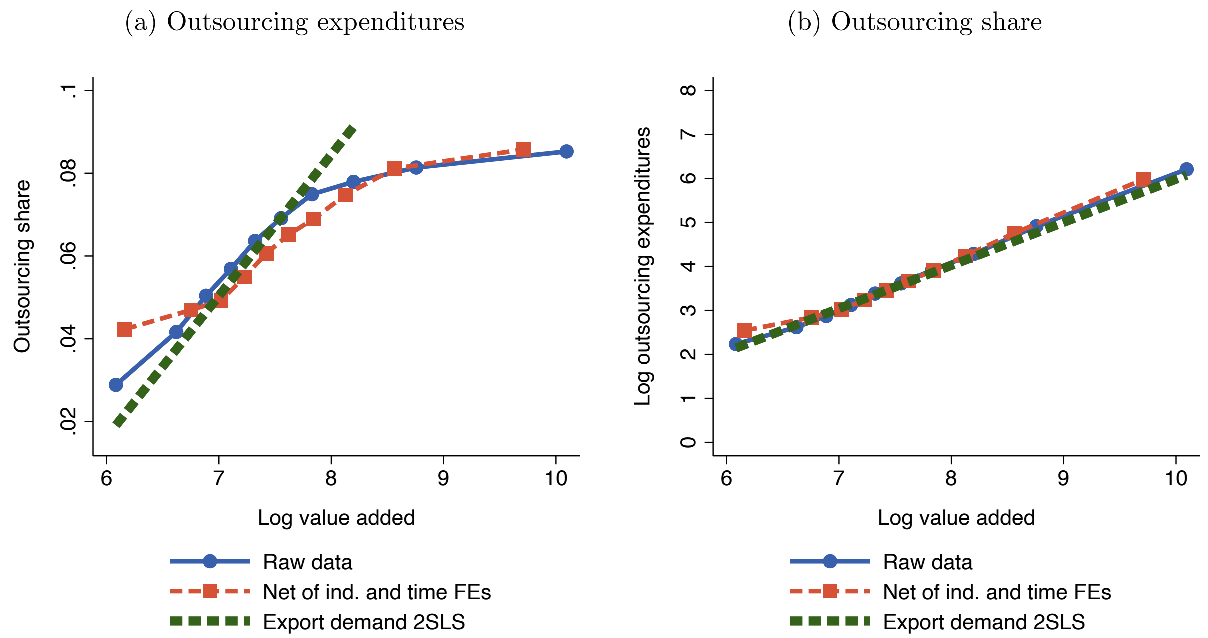

Note: Solid blue line : raw data. Dashed orange line: after removing 3-digit industry and time period fixed effects from the outsourcing share and log value added. Green line: 2SLS estimate using the export demand shift-share instrument in equation (11). Panel (a): log outsourcing expenditures. Panel (b): outsourcing share.

Figure 12: Components of export demand instrumental variable.

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

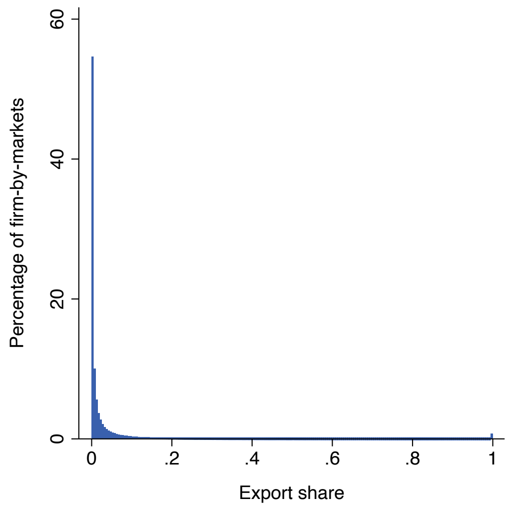

(b) Distribution of export demand by market.

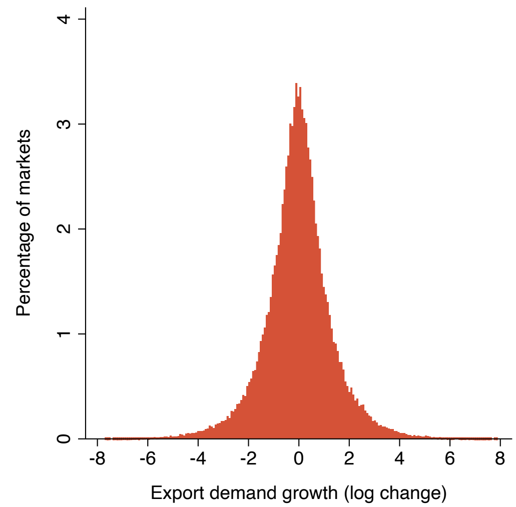

(c) Distribution of IV

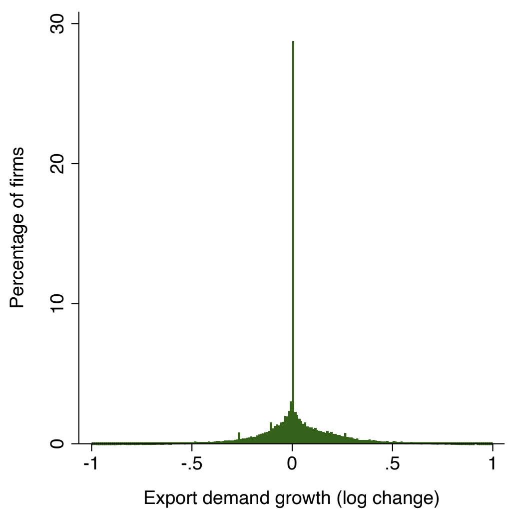

(d) Distribution of IV without zeros

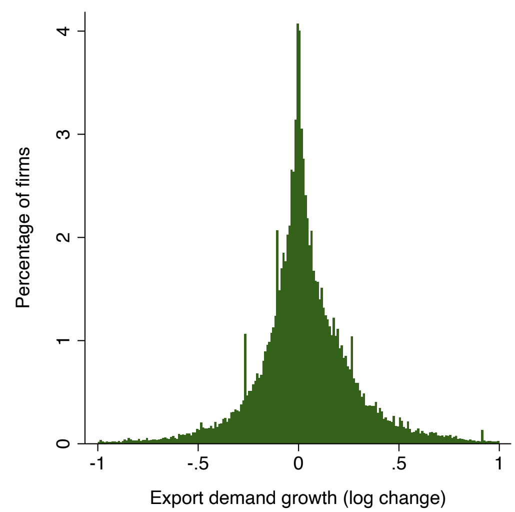

Note: Distributions of components of the instrumental variable Z_{ft} defined in equation (11). Panel (a) shows the distribution of the firm-by-market export shares $\pi_{f,t_0,j}$. Panel (b) shows that distribution of changes in export demand $\Delta \log X_{j,t,-f}$. Panel (c) shows the distribution of the IV Z_{ft}. Panel (d) shows the same distribution after removing zeros.

Figure 13: Selection into outsourcing: Reduced form of instrumental variable approach.

(a) Expenditures: reduced form

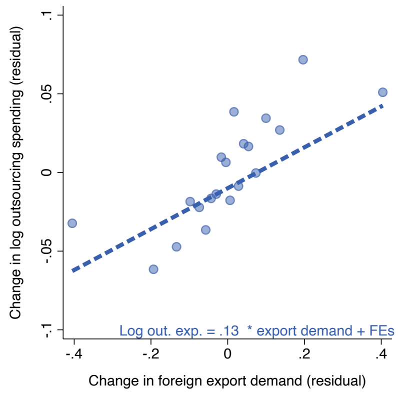

(b) Share: reduced form

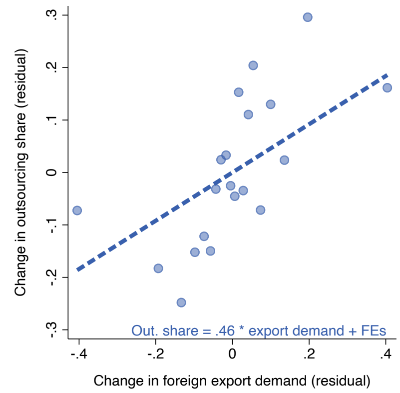

Note: Bin-scatterplot of the reduced form for selection into outsourcing with log outsourcing expenditures (panel (a)) and with the outsourcing share (panel (b)) as dependent variables. Panels include regression lines for 20 bins. Coefficients may differ from full sample regression coefficients reported in Tables 3 and 4, Online Appendix F.1.

Figure 14: Outsourcing by revenues.

(a) Outsourcing expenditures

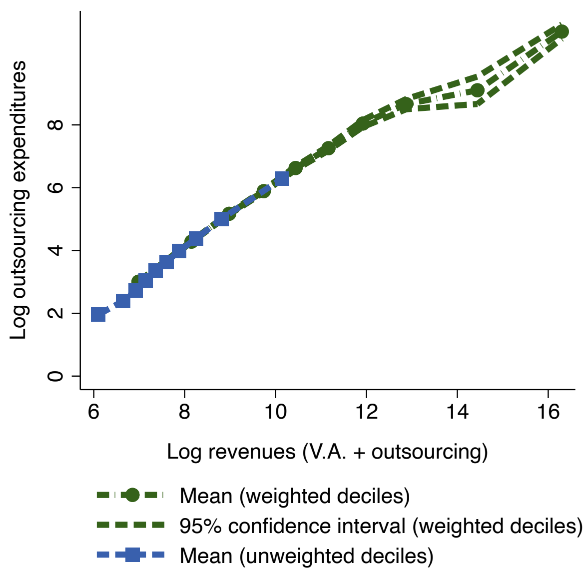

(b) Outsourcing share

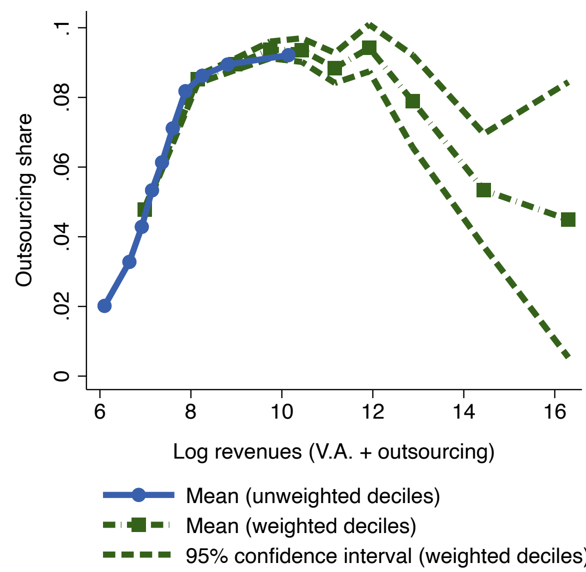

Note: Solid blue line : deciles based on number of firms. Dotted green line: deciles based on revenue-weighted quantiles. Dashed green lines: 95% confidence interval for revenue-weighted decile means. Panel (a): log outsourcing expenditures. Panel (b): outsourcing share.

F.2 The productivity effect

Figure 15: Components of outsourcing share instrumental variable.

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

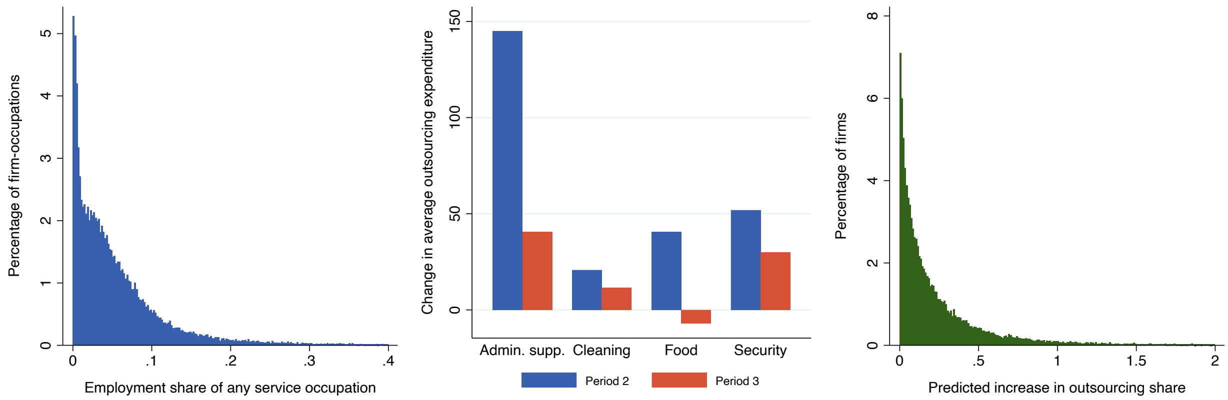

Note: Distributions of components of the instrumental variable Z'_{ft} for the firm-level outsourcing share. Panel (a) shows the distribution of the firm-by-occupation wage shares shares $\omega_{f,t_0,o}$ for service occupations. Panel (b) shows changes in average outsourcing expenditures by occupation $\Delta\Omega_{o,t,-f}$. 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.

Figure 16: Productivity effect using value added: Instrumental variable approach.

(a) First stage. (b) Reduced form. (c) 2SLS.

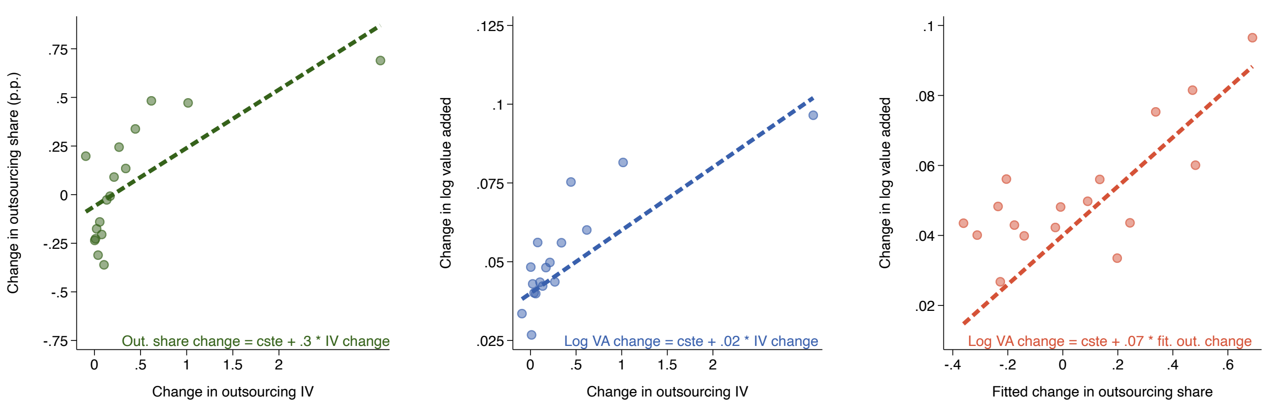

Note: Bin-scatterplot of the first stage (panel (a)), reduced form (panel (b)) and two stage least square (panel (c)) estimates for selection into outsourcing using value added VA_{ft} as a dependent variable. Panels include regression lines for 20 bins. Coefficients may differ from full sample regression coefficients reported in Table 5, Online Appendix F.2. Firm-level dataset (FICUS-EAE-DADS *Postes*-Trade).

F.3 The distributional effect

In this section we verify that larger firms pay service workers more and thus locate at the top of the job ladder. Following the large literature establishing 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:

$$\log w_{i,t} = \varphi_i + \psi_{I(i,t)} + \phi \log \text{Total employment}_{J(i,t)} + v_{i,t}. \tag{31}$$

i indexes workers, $I(i,t)$ is the industry of the employer of worker i in quarter t, $J(i,t)$ 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 10. 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.

Table 10: 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.

F.4 Alternative explanations

Figure 17 and Table 11 show outsourcing by industry volatility. Volatility calculated at the group (e.g. firm or industry) g level as the unweighted standard deviation of log value added growth across periods: $\text{vol}_{gt} = \sqrt{\mathbb{V}\text{ar}_{gt}[\log \text{VA}_{f,t+1} - \log \text{VA}_{ft}]}$.

Figure 17: Outsourcing by industry volatility.

(a) Volatility by value added and industry.

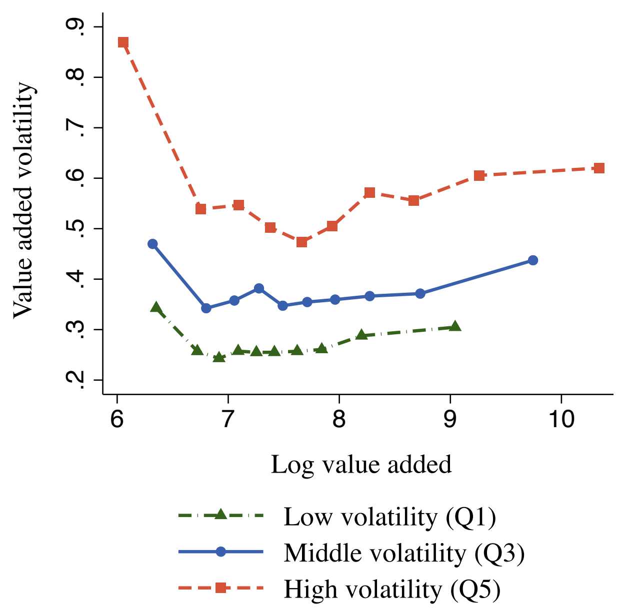

(b) Outsourcing by value added and industry.

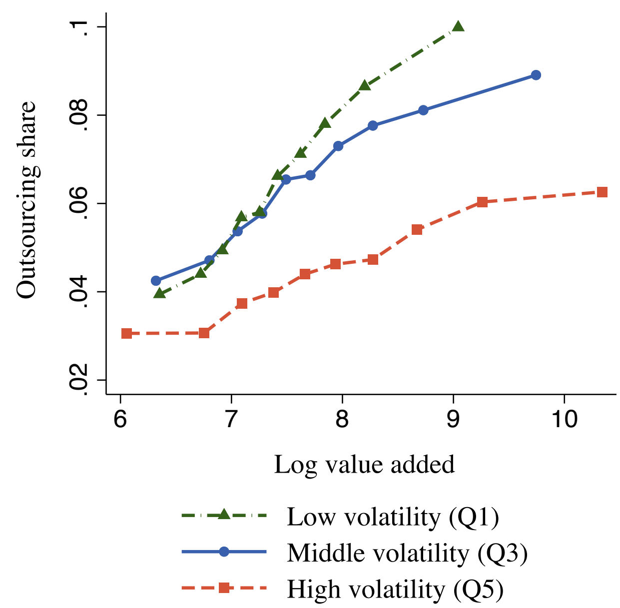

Figure 18: Size-dependent policies and outsourcing.

(a) Firm size distribution

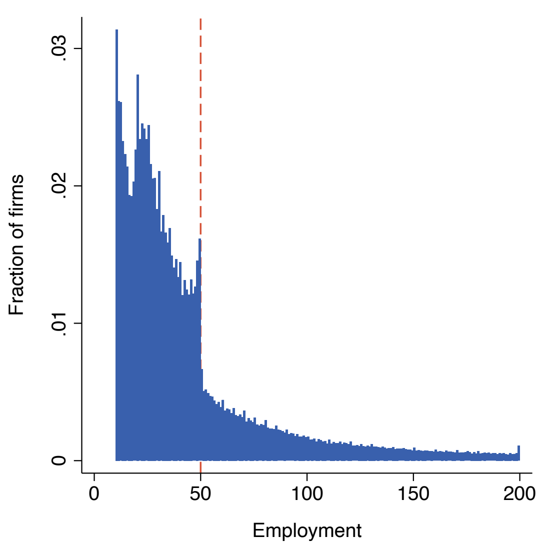

(b) Outsourcing by size

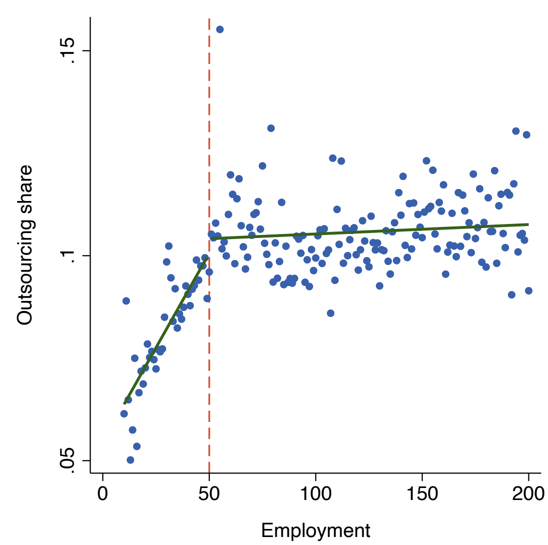

Table 11: Dependent variable: spending on external workers as a fraction of labor costs, in p.p.

	(1)	(2)	(3)	(4)	(5)	(6)
Log V.A.	1.991***		1.748***		1.858***	1.719***
	(0.017)		(0.111)		(0.023)	(0.105)
Volatility		0.605***		0.634***	0.624***	0.707***
		(0.022)		(0.036)	(0.022)	(0.037)
Fixed Effects						
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. Firm fixed effects excluded because firm-level volatility computed using all sample periods.

Table 12: 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***	-0.06***	-0.01**			
	(0.00)	(0.00)	(0.00)			
Wage P90-P10				-0.02***	-0.02***	-0.00*
				(0.00)	(0.00)	(0.00)
Log Revenues		0.02***	0.03***		0.02***	0.03***
		(0.00)	(0.00)		(0.00)	(0.00)
Fixed Effects						
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.

G Quantitative model

G.1 Model

Good producers Good producers solve the following problem

$$\pi(z,\varepsilon) = \max_{\boldsymbol{n},\boldsymbol{w},\boldsymbol{v},o} R(z,\boldsymbol{n}) - ((1-o)w_1 - op\varepsilon)\, n_1 - (1-o)c(v_1) - \sum_{s>1}(w_s n_s + c(v_s)), \tag{32}$$

subject to

$$n_s(w_s, v_s) = \frac{(1+k_s)e_s}{(1+k_s(1-F_s(w_s)))^2}\frac{v}{V_s} \quad \text{if } o_s = 0.$$

For future reference, we define $\eta_s(w_s) \equiv n_s(w_s, v_s)(V_s/v)$ as the number of workers per vacancy. 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 notation x^i and x^o to refer to the optimal choice of producers conditional on hiring low-skill workers in-house or via a contractor. The outsourcing decision is optimal: $o(z,\varepsilon) = 1\{\pi^o(z,\varepsilon) \geq \pi^i(z,\varepsilon)\}$.

We use the following parametric assumptions. The revenue function is Cobb-Douglas nested in a decreasing returns upper tier, $R(z,\boldsymbol{l}) = \left(z\prod_{s=1}^{S} l_s^{a_s}\right)^{\rho}$, with $\sum_s a_s = 1$. The vacancy cost function is isoelastic with elasticity γ, $c(v) = \mu_z\left(\frac{v^{1+\gamma}}{1+\gamma}\right)$, where we normalized the vacancy cost of a vacancy for good producers to the mean TFP, μ_z. The joint distribution of (z,ε), denoted by Γ, is log-normal:

$$(\log z, \log\varepsilon) \sim \mathcal{N}\left[\begin{pmatrix}\log\mu_z \\ 0\end{pmatrix}, \begin{pmatrix}\nu^2 & \iota \\ \iota & \sigma\end{pmatrix}\right].$$

We let $\boldsymbol{\Sigma}$ denote the variance-covariance matrix $\boldsymbol{\Sigma}$. As in Melitz (2003), there is an infinite supply of possible producers. Firms pay an entry cost η before drawing their productivity profile (z,ε). The mass of producers, M^g, is pinned down in expectation by the free entry condition $\mathbb{E}[\pi(z,\varepsilon)] = \eta$.

Contractors Contractors hire low-skill service workers to produce the outsourcing service through a constant returns to scale production function. They differ in their labor productivity, z. Their revenue function is $R^c(z,l) = p\tau zl$, where p is the price of the service and τ is a productivity wedge. Contractors' vacancy cost function shares the same elasticity as producers but with a different cost shifter: $c^c(v) = \mu_z\overline{c}^{\gamma}\left(\frac{v^{1+\gamma}}{1+\gamma}\right)$. Contractors solve the maximization problem

$$\pi^c(z) = \max_{n,w,v} R^c(z,n) - wn - c^c(v) \quad \text{s.t.} \quad n \leq n_1(w,v). \tag{33}$$

The distribution of z, denoted $\overline{\Psi}$, is log-normal with zero mean and standard deviation ν^c. As for producers, there is an infinite pool of potential contractors. Contractors pay an entry cost η^c before knowing their productivity z. The mass of contractors, $\overline{M}^c$, is determined by the free entry condition

$\mathbb{E}[\pi^c(z)] = \eta^c$.

Equilibrium conditions There are four sets of equilibrium conditions. The first define the wage offer distributions. For low-skill, service workers, both good producers hiring low-skill workers in-house and contractors compete on the job ladder. The wage offer distribution of skill 1 is[30]

$$F_1(w) = M^g \int\int 1\{w_1^i(z,\varepsilon) \le w\}(1 - o[z,\varepsilon]) \left(\frac{v_1^i(z,\varepsilon)}{\overline{V}_1}\right) d\Gamma(z,\varepsilon) + \\ \overline{M}^c \int 1\{\pi^c(z) \ge 0\} 1\{w_1^c(z) \le w\} \left(\frac{v_1^c(z)}{\overline{V}_1}\right) d\overline{\Psi}(z). \tag{34}$$

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 relative mass of vacancy attached to wages lower than w and offered by contractors. $\overline{V}_1$ is the aggregate mass of vacancy posted for low-skill workers:

$$\overline{V}_1 = M^g \int\int (1 - o[z,\varepsilon]) v_1^i(z,\varepsilon) d\Gamma(z,\varepsilon) + \overline{M}^c \int 1\{\pi^c(z) \ge 0\} v_1^c(z) d\overline{\Psi}(z). \tag{35}$$

For all the other skills $s > 1$, only good producers compete on the job ladder. However, the wage offered and vacancy posted for a skill $s > 1$ may differ across producers depending on whether low-skill workers are hired in-house or outsourced. Accordingly, the wage offer distribution of skill $s > 1$ is

$$F_s(w) = M^g \int\int 1\{w_s^i(z,\varepsilon) \le w\}(1 - o[z,\varepsilon]) \left(\frac{v_s^i(z,\varepsilon)}{\overline{V}_s}\right) d\Gamma(z,\varepsilon) + \\ M^g \int\int 1\{w_s^o(z,\varepsilon) \le w\} o(z,\varepsilon) \left(\frac{v_s^o(z,\varepsilon)}{\overline{V}_s}\right) d\Gamma(z,\varepsilon). \tag{36}$$

The first integral is the relative mass of vacancy attached to wages lower than w for skill s 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 for skill s offered by good producers outsourcing their service workers. Aggregate vacancy for skill $s > 1$ reads

$$\overline{V}_s = M^g \int\int (1 - o[z,\varepsilon]) v_s^i(z,\varepsilon) d\Gamma(z,\varepsilon) + M^g \int\int o(z,\varepsilon) v_s^o(z,\varepsilon) d\Gamma(z,\varepsilon), \tag{37}$$

The second set of equilibrium conditions determine the contact rates through the matching function:

$$\lambda_s^u = \mu_s \left(\frac{V_s}{u_s + \zeta_s(1 - u_s)}\right)^{1-\xi}, \tag{38}$$

[30]With decreasing returns to scale, producers always make positive profits. In contrast, low-productivity constraints, who use a constant returns to scale technology, may make negative profits.

where $u_s = 1/(1+k_s^U)$. The third equilibrium conditions pin down the reservation wages

$$\underline{w}_s = b_s + \left(k_s^U - k_s\right)\left(\int_{\underline{w}_s} \frac{1 - F_s(w)}{1 + k_s(1 - F_s(w))} \mathrm{d}w\right). \tag{39}$$

Finally, the last equilibrium condition requires the aggregate demand for the outsourcing service from good producers to equate the aggregate supply provided by contractors:

$$\overline{M}^c \int 1\{\pi^c(z) \geq 0\} \tau z n_1^c(z) d\overline{\Psi}(z) = M^g \int\int o(z,\varepsilon)\varepsilon n_1^o(z,\varepsilon) d\Gamma(z,\varepsilon), \tag{40}$$

where n_1^c and n_1^o are the numbers of low-skill workers hired by contractors and outsourcing producers respectively.

Equations (32) to (40) define jointly the equilibrium.

Definition 1 (Equilibrium). *An equilibrium is a collection of wage and vacancy functions for good producers, $\{w_s^\theta, v_s^\theta\}_{s\in\{1,\dots S\},\theta\in\{i,o\}}$, an outsourcing function for producers, o, wage and vacancy functions for contractor firms, w_1^C and v_1^C, wage distributions, $\{F_s\}_{s=1}^S$, aggregate vacancies, $\{\overline{V}_s\}_{s=1}^S$, reservation wages, $\{\underline{w}_s\}_{s=1}^S$, contact rates, $\{\lambda_s^u\}_{s=1}^S$, mass of producers and contractors, $\{\overline{M}^g, \overline{M}^c\}$, and outsourcing prices, p, such that*

1. *Given $\{F_s\}_{s=1}^S$, $\{\overline{V}_s\}_{s=1}^S$, $\{\lambda_s^u\}_{s=1}^S$, and p, the functions $\{w_s^\theta, v_s^\theta\}_{s=1}^S$ solve (32) conditional on an outsourcing decision;*
2. *Given $\{w_s^\theta, v_s^\theta\}_{s=1}^S$, $\{\lambda_s^u\}_{s=1}^S$, and p, the outsourcing decision is optimal $o(z,\varepsilon) = 1\{\pi^o(z,\varepsilon) \geq \pi^i(z,\varepsilon)\}$;*
3. *Given F_1, $\overline{V}_1$, λ_1^u, and p, the functions w_1^C and v_1^C solve (33);*
4. *Given $\{w_s^\theta, v_s^\theta\}_{s=1}^S$ and $\{w_1^c, v_1^c\}$, the wage distributions satisfy (34) and (36);*
5. *Given $\{v_s^\theta\}_{s=1}^S$ and v_1^c, aggregate vacancies are given by (35) and (37);*
6. *Given $\{\overline{V}_s\}_{s=1}^S$, the contact rates satisfy (38);*
7. *Given $\{F_s\}_{s=1}^S$ and $\{\lambda_s^u\}_{s=1}^S$, the reservation wages are given by (39);*
8. *Given $\{w_s^\theta, v_s^\theta, l_s^\theta\}_{s=1}^S$ and p, the mass of producers ensure zero profit in expectations;*
9. *Given $\{w_1^c, v_1^c, l_1^c\}$ and p, the mass of contractors ensure zero profit in expectations;*
10. *Given $\{w_1^c, v_1^c, l_1^o, o\}$ and p, the market clearing condition (40) holds.*

G.2 A tractable reformulation of the model

Due to its large amount of heterogeneity, this model is *a priori* numerically non-tractable. The standard fixed point of Burdett and Mortensen (1998) – that the wage and vacancy policy functions depend on the wage offer distributions, which themselves depend on the policy functions – is further complicated by the fact that multiple types of firms compete on the job ladder, and that the policy functions do not simply depend on firms' productivity.[31] In this section, we derive the optimality conditions behind the firms' problem (32) and (33). We then show how these can be rewritten to make the computation of the equilibrium numerically tractable. This tractability is feasible thanks to three assumptions: a Cobb-Douglas revenue function, a single outsourceable worker type, and a log-normal distribution for (z, ε).

Good producers: outsourced hires The problem of an outsourcing good producer reads

$$\pi^o(z, \varepsilon) = \max_{\boldsymbol{n}, \{w_s\}_{s>1}, \{v_s\}_{s>1}} R(z, \boldsymbol{n}) - n_1 p\varepsilon - \sum_{s>1} (w_s n_s + c(v_s)) \,. \tag{41}$$

Outsourcing good producers face an unconstrained hiring problem regarding how many low-skill service workers to hire. Taking the first-order condition yields

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

Plugged back into (41), the profit of the firm rewrites

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

for κ and $\mathcal{G}$ two parametric constants.[32] In the above expression, payroll and vacancy costs are independent from z and ε. Meanwhile, revenues only depend on the TFPR bundle $z\varepsilon^{a_1}$. As a result, the policy functions of type-o good producers are only a function of $z\varepsilon^{a_1}$ and it is not needed to keep track of z and ε separately. Since (z, ε) is jointly log-normally distributed, so is $(z, z\varepsilon^{a_1})$ and a closed-form expression exists for its variance-covariance matrix. Let Φ denote the log-normal distribution under the change of variable $(z, \varepsilon) \to (z, z\varepsilon^{a_1})$.

Good producers: in-house hires The previous derivation allows to treat the hiring problem of in-house and outsourcing good producers symmetrically. For both in-house and outsourcing firms, the

[31] In 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.

[32] Specifically, we have

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

hiring decisions of in-house workers follow two first order conditions:

$$\frac{\partial R^\theta[z,\boldsymbol{n}^\theta(z)]}{\partial n_s} - w_s^\theta(z) = \left(\frac{1+k_s[1-\Upsilon_s^\theta(z)]}{2k_s\partial_z\Upsilon_s^\theta(z)}\right)\partial_z w_s^\theta(z), \tag{42}$$

$$\left(\frac{\partial R^\theta[z,\boldsymbol{n}^\theta(z)]}{\partial n_s} - w_s^\theta(z)\right) n_s[w_s^\theta(z)] = \mu_z v_s^\theta(z)^{1+\gamma}, \tag{43}$$

where $\Upsilon_s^\theta(z) \equiv F_s[w_s^\theta(z)]$. For $\theta = i$, these two conditions have to hold for all $s \in \{1,\ldots,S\}$. For $\theta = o$, they have to hold for $s \in \{2,\ldots,S\}$.

Contractors The optimality condition of contractors are identical to that of producers. They are:

$$\frac{\partial R^c[z,n_1^c(z)]}{\partial n_1} - w_1^c(z) = \left(\frac{1+k_1[1-\Upsilon_1^c(z)]}{2k_1\partial_z\Upsilon_1^c(z)}\right)\partial_z w_1^c(z), \tag{44}$$

$$\left(\frac{\partial R^c[z,n_1^c(z)]}{\partial n_1} - w_1^c(z)\right) n_1[w_1^c(z)] = \mu_z[\bar{c}v_1^c(z)]^{1+\gamma}. \tag{45}$$

Plugging the vacancy optimality condition (45) into their profits function (33), it is easy to show that only contractors with productivity $z \geq \underline{w}_1/p\tau \equiv \underline{z}^c$ make positive profits.[33]

As in the standard Burdett and Mortensen (1998), equation (42) is a differential equation that can be solved forward subject to some initial condition on $w_s^\theta(\underline{z})$. There are, however, three differences with the standard Burdett and Mortensen (1998). First, due to the presence of decreasing returns to scale, the MPL of each worker hired in-house depends on the number of workers hired of other skills. Accordingly, (43) must be solve jointly for every skill hired in-house. Second, good producers hiring service workers in-house compete with contractors on the job ladder through (34). Solving for (42) therefore require to know the policy functions of contractors given by (44) and (45). Third, good producers hiring service workers in-house compete through (36) for the other skills with good producer outsourcing service workers. As such, solving for (42) requires to know the policy functions of both in-house and outsourcing firms.

These three differences altogether complicate the numerical fixed point to be solved for. We now show how the existence of three functions allow us to simplify the computation of (42) by only solving the differential equation for in-house producers. The first function characterizes the outsourcing decision. The second and third functions map the wages of contractors and outsourcing producers respectively on that of in-house producers.

Outsourcing choice The first function is the indifference function between hiring service workers in-house versus outsourcing the hiring. This function simplifies the computation of the wage distribution, (34) and (36). 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 $z\varepsilon^{-a_1} > \varphi(z)$. Given

[33] As in the original Burdett and Mortensen (1998), the CRS technology ensures that the least productive contractors, $\underline{z}^c$, pays the reservation wage to its workers. The marginal revenues of a hire are $p\tau\underline{z}^c$.

φ, the wage offer distribution for service workers, (34), is given by

$$F_1(w) = \frac{M^G}{V_1} \int \mathbf{1}\{w_1^i(z) \leq w\} v_1^i(z) \Omega^i(z) \mathrm{d}\Phi_z(z) + \frac{\overline{M}^C}{V_1} \int_{\underline{w}/p\tau} \mathbf{1}\{w_1^c(z) \leq w\} v_1^c(z) \mathrm{d}\overline{\Psi}_Z(z).$$

The function $\Omega^i(x)$ is the probability that a good producer with productivity z hires its service workers in-house, $\Omega^i(x) \equiv \Phi_{\hat{Z}|Z}[\varphi(x) \mid x]$. Accordingly, from the definition of Υ_1^i, we have

$$\Upsilon_1^i(z) = \frac{M^G}{V_1} \int^z v_1^i(\zeta) \Omega^i(\zeta) \mathrm{d}\Phi_z(\zeta) + \frac{\overline{M}^C}{V_1} \int_{\underline{w}/p\tau} \mathbf{1}\{w_1^c(\zeta) \leq w_1^i(z)\} v_1^c(\zeta) \mathrm{d}\overline{\Psi}_Z(\zeta). \tag{46}$$

For skill $s > 1$, the wage offer distribution (36) rewrites

$$F_s(w) = \frac{M^G}{V_s} \left(\int \mathbf{1}\{w_s^i(z) \leq w\} v_s^i(z) \Omega^i(z) \mathrm{d}\Phi_Z(z) + \int \mathbf{1}\{w_s^o(\hat{z}) \leq w\} v_s^o(\hat{z}) \Omega^o(\hat{z}) \mathrm{d}\Phi_{\hat{Z}}(\hat{z}) \right).$$

The function $\Omega^o(x) \equiv \Phi_{Z|\hat{Z}}[\varphi^{-1}(x) \mid x]$ is the probability that a good producer with TFPR bundle x outources its service workers. As for service workers, we have

$$\Upsilon_s^i(z) = \frac{M^G}{V_s} \left(\int^z v_s^i(\zeta) \Omega^i(\zeta) \mathrm{d}\Phi_Z(\zeta) + \int \mathbf{1}\{w_s^o(\hat{z}) \leq w_1^i(z)\} v_s^o(\hat{z}) \Omega^o(\hat{z}) \mathrm{d}\Phi_{\hat{Z}}(\hat{z}) \right). \tag{47}$$

The next two functions simplify the computation of the second integral in (46) and (47) by allowing us to recover the policy functions of contractors and outsourcing good producers from that of in-house good producers.

Contractors vs. producers Let $\zeta^{G\to C}$ be the productivity of a contractor that posts the same wage as a good producer with productivity z; that is, $w_1^i(z) = w_1^C[\zeta^{G\to C}(z)]$. Since $w_1^i(\underline{z}) \geq \underline{w}_1 = w_1^c(\underline{z}^c)$ and $\lim_{z\to\infty} w_1^i(z) = \overline{w}_1 = \lim_{z\to\infty} w_1^c(z)$, the function $\zeta^{G\to C}$ is defined on the support of producers' productivity.

A closed-form expression exists for $\zeta^{G\to C}$. For all $z \geq \underline{z}$, $w_1^i(z) = w_1^C[\zeta^{G\to C}(z)]$ implies by definition that these firms have the same rank in the wage offer distribution. In particular, the right-hand side of (44) and (42) must be equal, and therefore so must be the marginal product of labor of good producer z and contractor firms $\zeta^{G\to C}(z)$. Since the left-hand side of (44) and (42) are equal, so must be the right-hand side of (45) and (43), which implies

$$v_1^c[\zeta^{G\to C}(z)] = \frac{v_1^i(z)}{\overline{c}}.$$

This equivalence allows us to then invert the MPL equality condition to recover an expression for

$\zeta^{G\to C}$:

$$\zeta^{G\to C}(z) = \frac{1}{p\tau}\frac{\partial R^i[z, \boldsymbol{n}^i(z)]}{\partial n_1} \tag{48}$$

Going back to the wage distribution (46), we can differentiate it with respect to z to obtain

$$d\Upsilon_1^i(z) = \frac{M^G}{V_1} v_1^i(z)\Omega^i(z)\mathrm{d}\Phi_z(z) + \frac{\overline{M}^C}{V_1}\frac{v_1^i(z)}{\bar{c}}\mathrm{d}\overline{\Psi}_Z[\zeta^{G\to C}(z)]. \tag{49}$$

Ignoring for now the complementarity across skills, equations (42) and (49) constitute a system of differential equations. Combined with the vacancy optimality condition, (43), and two boundary conditions, $w_1^i(\underline{z}) = \underline{w}_1^i$ and $\Upsilon_1^i(\underline{z}) = \underline{F}_1^i$, these differential equations yield the policy functions $(w_1^i, v_1^i, w_1^c, v_1^c, F_1)$ without solving for the policy functions of contractors.

Absent contractors on service workers' job ladder or decreasing returns to scale, the boundary conditions would be $\underline{w}_1^i = \underline{w}_1$ and $\underline{F}_1^i = 0$. However, the presence of both features imply that it may be optimal for producers *not* to post the reservation wage, $\underline{w}_1^i > \underline{w}_1$ and $\underline{F}_1^i > 0$. When $\underline{w}_1^i > \underline{w}_1$, all wages between $[\underline{w}_1, \underline{w}_1^i)$ are offered by contractor firms, which are therefore alone on the job ladder. The policy functions of contractors on $[\underline{w}_1, \underline{w}_1^i)$ are similar to that of a standard vacancy-extended Burdett and Mortensen (1998) problem:

$$\begin{aligned} \partial_z w_1^c(z) &= \left(\frac{2k_s\partial_z\Upsilon_1^c(z)}{1 + k_s[1 - \Upsilon_1^c(z)]}\right)[p\tau z - w_1^c(z)]\,, \\ d\Upsilon_1^c(z) &= \frac{\overline{M}^C}{V_1}\frac{1}{\bar{c}}\left[(p\tau z - w)\,\frac{\eta_1(w)}{\mu_z V_1}\right]^{\frac{1}{\gamma}} d\overline{\Psi}_Z(z), \end{aligned} \tag{50}$$

where Υ_1^c is defined in an analoguous fashion as Υ_1^i. The two differential equations are subject to the boundary conditions $w_1^c(\underline{z}) = \underline{w}_1$ and $\Upsilon_1^c(\underline{z}) = 0$. The system (50) pins down wages and the wage offer distribution for contractors with TFP in $[\underline{z}^C, \tilde{z}^C)$, where $\tilde{z}^C$ is such that $w_1^C(\tilde{z}^C) = \underline{w}_1^i$. By a similar argument as before, this wage equality implies that the MPL of the two firms have to be equal:

$$\tilde{z}^C = \frac{1}{p\tau}\frac{\partial R^i[\underline{z}, \boldsymbol{n}^i(\underline{z})]}{\partial n_1}. \tag{51}$$

In-house vs. outsourced producers A similar derivation exists to obtain an equivalence between the policy functions of in-house versus outsourced producers. Let $\zeta_s^{i\to o}$ be the productivity (bundle) of an outsourced producer that posts the same wage as an in-house producer with productivity z for skill s; that is, $w_s^i(z) = w_s^o[\zeta^{i\to o}(z)]$. As for contractors, two producers that offer the same wage have the same MPL for those workers and decide to post the same number of vacancies. These firms therefore attain the same size, and the following equality holds:

$$\frac{R^i[z, \boldsymbol{l}^i(z)])}{n_s^i(z)} = \left(\frac{\kappa}{\rho}\right)\left(\frac{R^o[\zeta_s^{i\to o}(z), \boldsymbol{n}^o(\zeta_s^{i\to o}(z))]}{n_s^o[\zeta_s^{i\to o}(z)]}\right) = \left(\frac{\kappa}{\rho}\right)\left(\frac{R^o[\zeta_s^{i\to o}(z), \boldsymbol{n}^i(z)]}{n_s^i(z)}\right). \tag{52}$$

Said differently, it must be that $R^o[\zeta_s^{i\to o}(z), \boldsymbol{n}^i(z)] = \rho R^i[z, \boldsymbol{n}^i(z)]/\kappa$ for all s, and the function $\zeta_s^{i\to o}(z) \equiv \zeta^{i\to o}(z)$ is skill-independent. Inverting equation (52) yields a closed-form expression for $\zeta^{i\to o}$:

$$\zeta^{i\to o}(z) = p^{a_1}\left(\frac{\rho}{\kappa\mathcal{G}}\right)^{\frac{1}{\kappa}}\left(zn_1^i(z)^{a_1}\prod_{s>1} n_s^i(z)^{a_s\left(1-\frac{\kappa}{\rho}\right)}\right)^{\frac{\rho}{\kappa}}. \tag{53}$$

Finally, this function can be used in (47) to obtain the differential equation:

$$d\Upsilon_s^i(z) = M^G\left(\frac{v_s^i(z)}{V_s}\right)\left(\Omega^i(z)\mathrm{d}\Phi_Z(z) + \Omega^o[\zeta^{i\to o}(z)]\mathrm{d}\Phi_{\hat{Z}}[\zeta^{i\to o}(z)]\right). \tag{54}$$

Here as well, ignoring for now the skill complementarity, (42) and (54), together with the vacancy optimality condition (43), form a system of differential equations which, subject to the boundary conditions $\underline{w}_s^i(\underline{z}) = \underline{w}_s$ and $\Upsilon_s^i(\underline{z}) = 0$, returns the policy functions (w_s^i, v_s^i) and the job ladder Υ_s^i.[34]

G.3 Algorithm

The algorithm has four levels of iteration. The most inner level solves for the functions, $\{w_s^\theta, v_s^\theta, \Upsilon_s^\theta\}_{s,\theta}$ using the system of differential equations obtained in Section G.2. The second most inner levels iterate on the aggregate number of vacancy, $\{V_s\}_s$, and search frictions, $\{\lambda_s^U, \lambda_s^E\}$. The intermediate levels iterate on the outsourcing decision, φ, and the reservation wages, $\{\underline{w}_s\}_s$. Finally, the outer level iterates on the market clearing and free entry conditions to solve for the price of outsourcing p and the mass of firms $\{\overline{M}^c, M^g\}$.

Differential equations Given an outsourcing price p, the indifference function φ, the reservation wages $\{\underline{w}_s\}_s$, the aggregate vacancies $\{V_s\}_s$, the contact rates $\{\lambda_s^U, \lambda_s^E\}_s$, and the mass of firms $\{\overline{M}^c, M^g\}$, the most-inner loop iterates forward on the differential equations to solve for the functions $\{w_s^\theta, v_s^\theta, \Upsilon_s^\theta\}_{s,\theta}$. 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.

We first iterate on the differential equations for the good producers that hire their service workers in-house. This requires to know the boundary conditions $\underline{w}_1^i$ and $\underline{F}_1^i$. 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 $s > 1$. The labor supplies are $\eta_1^i(\underline{z}) = \eta_1[w_1^c(z)]$ and $\eta_s^i(\underline{z}) = \eta_s(\underline{w}_s)$ for $s > 1$. Vacancies are solved jointly by iterating on (43).

[34] To avoid combinatorial issues, we assume that the least productive good producers post the reservation wage regardless of their outsourcing decision, $w_s^i(\underline{z}) = w_s^o(\underline{z}) = \underline{w}_s$ for $s > 1$. We verify numerically ex-post whether a profitable deviation exists.

2. Check if condition (51) holds at z. If it does, then $\underline{w}_1^i = w_1^c(z)$ and $\underline{F}_1^i = \Upsilon_1^c(z)$. If not, compute $\partial_z w_1^C(z)$ and $\partial_z \Upsilon_1^C(z)$ from (50) and go back to step 1.

Then, for each $z \geq \underline{z}$ starting from $\underline{z}$ for which we know that $w_1^i(\underline{z}) = \underline{w}_1^i$, $\Upsilon_1^i(\underline{z}) = \underline{F}_1^i$, and $w_s^i(\underline{z}) = \underline{w}_s$ and $\Upsilon_s^i(\underline{z}) = 0$ for $s > 1$:

1. Compute $\{\eta_s^i(z)\}_s$ from the labor supplies.

2. Compute $\{v_s^i(z)\}_s$ jointly by solving (43).

3. Compute $\zeta^{G\to C}(z)$ from (48) and $\zeta^{i\to o}(z)$ from (53). Compute also $\zeta^{G\to C}(z')$ and $\zeta^{i\to o}(z')$, where $z' = z + \Delta_z$ is the next point on the grid of z and we extrapolate variables linearly.

4. Use (49) to compute $\Upsilon_1^i(z')$.

5. Compute wages $\{w_s^i(z')\}_s$ from the wage ODEs (42).

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. The process is similar as above except for step 4: in this iteration, it is not required to compute $\zeta^{G\to C}$ but it is needed to find the numerical inverse of $\zeta^{i\to o}$. The functions $\{w_s^\theta, v_s^\theta, \Upsilon_s^\theta\}_{s,\theta}$ are therefore computed using only two iterations which allows for a fast computation of the equilibrium despite its complexity.

Inner loop Given an outsourcing price p, the function φ, the reservation wages $\{\underline{w}_s\}_s$, and the mass of firms $\{\overline{M}^c, M^g\}$ the inner loop iterates over the aggregate vacancies $\{V_s\}_s$ through (35) and (37) and the contact rates $\{\lambda_s^u\}_s$ through (38).

Intermediate loop Given an outsourcing price p and the mass of firms $\{\overline{M}^c, M^g\}$, the intermediate loop solves for the indifference function, φ, and the reservation wages, $\{\underline{w}_s\}_s$. 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 (39).

Outer loop The outer iteration solves for the price of outsourcing through the market clearing condition (40) and the mass of firms $\{\overline{M}^c, M^g\}$ through the free entry conditions $\mathbb{E}[\pi(z,\varepsilon)] = \eta$ and $\mathbb{E}[\pi^c(z)] = \eta^c$.

G.4 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 describe 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 i thus 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,

$$\begin{aligned}
\text{Ag. output} &= \sum_i \text{VA}_i^G + \sum_j \text{VA}_j^C \\
&= \sum_i \left(R_i^G - p\varepsilon_i n_i^G\right) + p\sum_j \tau z_j n_j^C \\
&= \sum_i \left(R_i^G - p\varepsilon_i n_i^G\right) + p\sum_i \varepsilon_i n_i^G \\
&= \sum_i R_i^G
\end{aligned} \tag{55}$$

where the first equality uses the definitions of value added, and the second equality uses labor services market clearing (40).

Aggregate TFP We define TFP as

$$\text{TFP} = \frac{\text{Ag. output}}{\overline{N}}, \tag{56}$$

where the labor aggregator $\overline{N}$ is defined as

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

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

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

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

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

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

$$\text{TFP} = \frac{\text{Ag. output}}{\widetilde{N}} \times \frac{\widetilde{N}}{\overline{\overline{N}}},$$

and so changes in aggregate TFP are

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

H Estimation

We estimate the quantitative model in three steps. In the first step, we use a benchmark estimate from the literature to calibrate the matching function elaslticity, $\chi = 0.5$. In the second step, we invert the model to read off our administrative data most of the parameters without the need for any simulation (about 80% of the parameters). In the third step, we estimate the remaining parameters by indirect inference. Section H.1 goes over the data used in the estimation. Section H.2 describes the various steps of the model inversion. Finally, Section H.3 presents the moments used in the indirect inference.

The estimated parameters are presented in Table 1.

H.1 Data

We combine three datasets: Ficus, the EAE, and the DADS Panel. Ficus provides information on firms' value added. The EAE gives us information on firms' outsourcing expenditures. Finally, the DADS Panel has information on firms' employment and wages per skill. The three datasets are merged on firms' ID.

We group firms in the data into three categories: contractors as identified by their 3-digit industry code, producers who outsource, and producers who do not outsource. Within each category, we group firms into 100 bins of average unconditional wage to reduce measurement error and focus on the heterogeneity relevant to our analysis. This grouping is consistent with our model in which more productive firms pay higher wages. It allows us to project the data on the relevant dimension of heterogeneity.

We do the following adjustment on the data. First, we make the data consistent with the model. We suppose that contractors only hire service workers. We assume that outsourcing firms do not hire any service workers. We set outsourcing expenditures to zero for in-house firms. For each adjustment, we update firms' value added by treating these expenditures as intermediaries. We then scale gross wages so as to obtain an empirically-relevant labor share of 70%.

Second, we account for worker selection between contractors and producers by setting the average wage gap between the two types of firms to the AKM gap computed in Section 3.3.

Third and last, we ensure that the three datasets are consistent with each other. In particular, FICUS possesses information on firms' aggregate wage bill. To guarantee that the value added in Ficus is consistent with the wage and size data in the DADS Panel, we project firms' value added on their wage bill in FICUS. We then use this relationship to compute the VA consistent with the wage bill in the DADS Panel. We proceed in a similar fashion to render outsourcing expenditures consistent with the information contained in the DADS Panel.

The scale of prices is defined up to a constant. We therefore normalize value added, outsourcing expenditures, and wages, by the average wage of the economy.

H.2 Model inversion

Labor market frictions The job destruction rate is identified from employment to non-employment transitions: $\delta_s = \mathrm{EN}_s$. The contact rate of unemployment maps into the rate of non-employment to employment transitions: $\lambda_s^U = \mathrm{NE}_s$. These pin down implicitly the matching function efficiency through (38).

We map the job-to-job transition rate in the model to the EE transition rate. Omit s indices for simplicity. Our argument requires only that the economy be stationary. Index firms by their wage offer w and the 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:

$$\text{EE}_s = \frac{\delta_s}{k_s}\left[(1+k_s)\log(1+k_s) - k_s\right].$$

This identifies ζ_s given δ_s.

Skills We observe the skill employment share, $e_s/\sum_{s'} e_{s'}$. In the model, the mass of employed workers of a particular skill is related to the total mass of workers of that skill through $e_s = k_s/(1+k_s)m_s$. Hence,

$$\frac{e_s}{\sum_{s'} e_{s'}} = \frac{k_s/(1+k_s)m_s}{\sum_{s'} k_{s'}/(1+k_{s'})m_{s'}} \quad \text{and} \quad \sum_s m_s = 1$$

jointly identifies $\{m_s\}_s$.

UI In the model, the replacement rate is $\text{RE}_s = b_s/\mathbb{E}_s[w]$. We set the non-employment insurance b_s to guarantee a replacement rate of 40%.

Vacancy and MPL We use the wage and size data to infer firms' vacancy share and marginal product of labor (MPL) for each skill hired in house. Specifically, fix a skill s, and take a firm j of type $\theta \in \{c, i, o\}$ in the sample. Then:

1. Compute the rank of the firm in the employment-weighted wage distribution, G^θ_{js}.
2. Compute the rank of the firm in the wage offer distribution:

$$F^\theta_{js} = \frac{G^\theta_{js}(1+k_s)}{1+k_s G^\theta_{js}}.$$

3. Compute the firm's size per vacancy posted from the labor supply curve:

$$\eta^\theta_{js} = \frac{(1+k_s)e_s}{[1+k_s\overline{F}^\theta_{js}]^2}.$$

Figure 19: Wage offer distributions

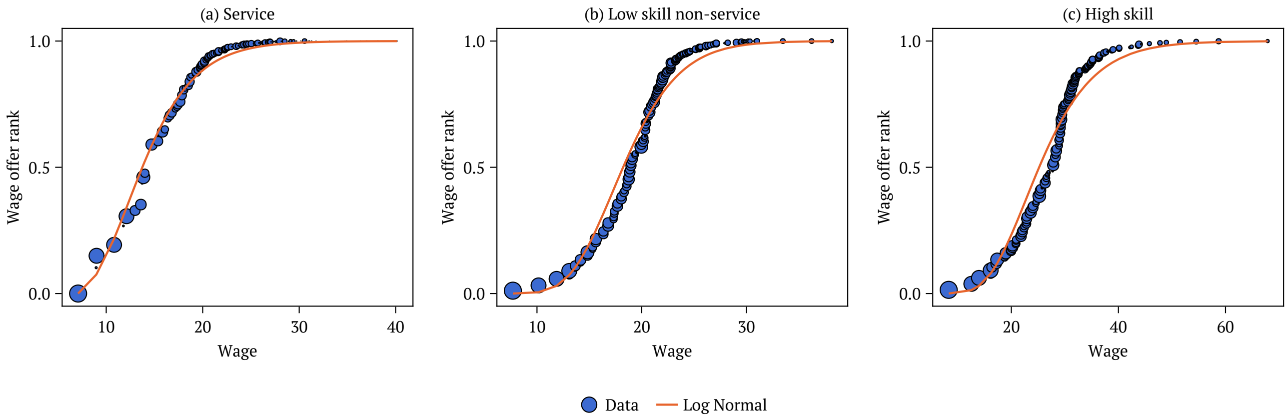

4. Compute the firm's vacancy share from its total size: $v^{\theta}_{js}/\overline{V}_s = n^{\theta}_{js}/\eta^{\theta}_{js}$.

5. Invert the wage optimality condition (42) to identify firms' MPL:

$$\mathrm{MPL}^{\theta}_{js} = w^{\theta}_{js} + \left(\frac{1 + k_s(1 - F^{\theta}_{js})}{2k_s}\right) \frac{\partial w}{\partial F}\bigg|_{w=w^{\theta}_{js}}. \tag{57}$$

To compute the partial derivative on the right-hand side of (57) while minimizing measurement error, we fit a log-normal distribution on the wage offer distribution F_s. We obtain the wage derivative as the inverse of the pdf, $\partial w/\partial F = 1/f_s(w)$ for $w = F_s^{-1}(F^{\theta}_{js})$. To further reduce measurement error, we winsorized the within-skill MPL distribution at the 2.5%. Figure 19 shows the winzorized wage offer distributions and the log-normal fit.

Mass of firms We identify the mass of producers and contractors from two equilibrium conditions. First, consistency requires that the aggregate demand for workers equate the aggregate supply for each skill. This implies

$$(M^c + M^g)\,\mathbb{E}[n_1] + M^g \sum_{s>1} \mathbb{E}[n_s] = \sum_s e_s. \tag{58}$$

Second, the outsourcing market clearing condition (40) requires the aggregate outsourcing expenditures to equate contractors' aggregate revenues. Expressed in terms of averages, this condition reads

$$\frac{M^c}{M^g} = \frac{\mathbb{E}[p\varepsilon_j n_{1j} 1\{o_j = 1\}]}{\mathbb{E}[p\tau z^c_j n^c_{1j}]}. \tag{59}$$

The numerator is the average outsourcing spending of producers. The denominator is the average revenue of contractors, which can be expressed as $p\tau z^c_j n^c_{1j} = \mathrm{MPL}^c_{j1} n^c_{1j}$ thanks to the CRS assumption. Accordingly, the right-hand side of (59) is measurable. Equations (58) and (59) together pin down

$\{M^c, M^g\}$.[35]

Producers' revenue function The parameters of the revenue function for good producers, ρ and $\{a_s\}_s$, can be identified from firms' MPL. For low-skill service workers hired in-house, their MPL is such that

$$\frac{\text{MPL}^i_{jl}}{R^i_j/n^i_{lj}} = \rho a_l,$$

where R^i_j is firm j's value added. Meanwhile, for firms outsourcing their low-skill service workers, we have

$$\frac{\text{OS}^o_j}{R^o_j} = \left(\frac{\rho}{1-\rho a_l}\right) a_l,$$

for OS_j firm j's outsourcing expenditures. Their weighted average identifies service workers' effective productivity,

$$\left(\frac{M^i}{M^g}\right)\mathbb{E}\left[\frac{\text{MPL}^i_{jl}}{R^i_j/n^i_{jl}}\right] + \left(1-\frac{M^i}{M^g}\right)\mathbb{E}\left[\frac{\text{OS}^o_j}{R^o_j}\right] = \left(\frac{M^i}{M^g}\right)\rho a_l + \left(1-\frac{M^i}{M^g}\right)\left(\frac{\rho a_l}{1-\rho a_l}\right),$$

where M^i/M^g is the fraction of producers hiring service workers in-house. Likewise, for the other skill $s > 1$,

$$\frac{\text{MPL}^i_{js}}{R^i_j/n^i_{js}} = \rho a_s \quad \text{and} \quad \frac{\text{MPL}^o_{js}}{R^o_j/n^o_{js}} = \left(\frac{\rho}{1-\rho a_l}\right) a_s,$$

such that

$$\mathbb{E}\left[\frac{\text{MPL}^\theta_{js}}{R^\theta_j/n^\theta_{js}}\right] = \left\{\left(\frac{M^i}{M^g}\right) + \left(1-\frac{M^i}{M^g}\right)\left(\frac{1}{1-\rho a_l}\right)\right\}\rho a_s.$$

The above expression identifies ρa_s. The normalization $\sum_s a_s = 1$ separates $\{a_s\}$ from ρ.

Producers' produtivity To identify producers' productivity, recall that value added is

$$\text{VA}^i_j = \left(z_j \prod_s n^{a_s}_{js}\right)^\rho \quad \text{and} \quad \text{VA}^o_j = \mathcal{G}\left(z(p\varepsilon)^{-a_l}\prod_{s\neq l} n^{a_s}_{js}\right)^\kappa$$

[35]When performing the indirect inference step, we ensure that (59) holds in the model.

Figure 20: Revenue TFP and wages

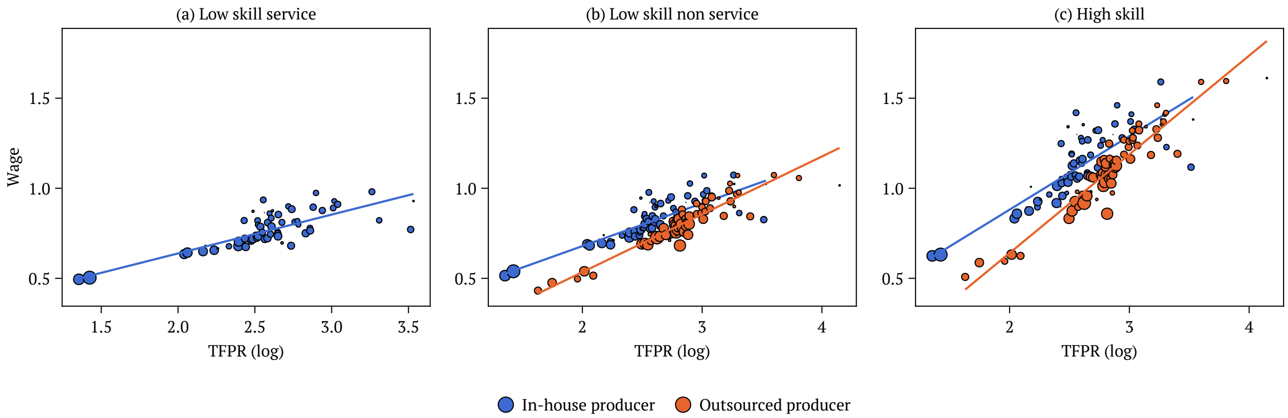

for in-house and outsourced producers respectively. Given our estimates of ρ and $\{a_s\}_s$, we can invert those two expressions to obtain

$$z_j = \frac{[\mathrm{VA}_j^i]^{1/\rho}}{\prod_s n_{js}^{a_s}} \quad \text{and} \quad z_j(p\varepsilon_j)^{-a_l} = \frac{[\mathrm{VA}_j^o/\mathcal{G}]^{1/\rho}}{\prod_{s\neq l} n_{js}^{a_s}}. \tag{60}$$

These identify the dispersion of log TFP for in-house and outsourced producers. Figure 20 displays wages against firms' productivity, confirming that wages are empirically increasing in productivity.

Vacancy cost elasticity Taking logs, the vacancy optimally condition of producers (43) reads

$$\log \frac{v_{js}^\theta}{V_s} = \alpha_s + \frac{1}{\gamma} \log \left[\eta_{js}^\theta \left(\mathrm{MPL}_{js}^\theta - w_{js}^\theta\right)\right],$$

where α_s is a skill-specific constant. The vacancy optimality condition therefore estimates the elasticity of the vacancy cost function:

$$\gamma = \sqrt{\frac{\sum_s e_s \mathbb{V}\mathrm{ar}_s \left[\log \left(\eta_{js}^\theta \left(\mathrm{MPL}_{js}^\theta - w_{js}^\theta\right)\right)\right]}{\sum_s e_s \mathbb{V}\mathrm{ar}_s \left[\log v_{js}^\theta / V_s\right]}}. \tag{61}$$

Contractors' productivity Contractors' MPL reads $\zeta_j \equiv p\tau z_j$ for each firm active in the economy. Given our parametric assumption on contractors' TFP, the distribution of MPL is log-normal with mean $\log \mu^c \equiv \log p\tau$ and standard deviation ν^z. We observe its left-truncated distribution at $\underline{z}^c$ – which we observe. We therefore jointly estimate μ^c and σ^c by method of moments. Finally, given $\{\mu^c, \nu^c\}$, we can compute the unconditional mass of contractors $\overline{M}^c$ from

$$\overline{M}^c = \frac{M^c}{1 - \overline{\Psi}_Z\left(\underline{z}^c\right)}.$$

Contractors' vacancy cost Given γ, contractors' relative vacancy cost is identified from their optimality condition, (45):

$$\log \bar{c} = \mathbb{E}\left[\hat{\beta} \log\left(\text{MPL}_{j1}^c - w_{j1}^c\right) \eta_{j1}^c + \hat{\alpha}_l - \log \frac{v_{j1}^c}{V_1}\right]. \tag{62}$$

Entry cost Producers' and contractors' entry cost are set to ensure zero expected profits:

$$\mathbb{E}[\pi(z, \varepsilon)] = \eta \quad \text{and} \mathbb{E}[\pi^c(z)] = \eta^c.$$

In practice, expected profits are computed from the simulated model to ensure consistency.

H.3 Indirect inference

Given the model inversion, there remains to estimate the parameters that dictate the productivity distribution of producers, $\{\mu_z, \mathbf{\Sigma}\}$, and contractors' productivity wedge, τ. We calibrate these moments jointly by indirect inference. We set the average TFP to ensure that the average wage in the model is one.[36] We calibrate producers' TFP dispersion, ν, to match the standard deviation of in-house producers' TFP as identified by (60). Absent selection into outsourcing, that moment exactly identifies ν. We set the dispersion in outsourcing cost, σ, to match the standard deviation in outsourced producers' TFP. Here as well, conditional on ν and ι and absent selection, that moment identifies σ. We calibrate the covariance between TFP and outsourcing cost to match the (standardized) value added elasticity of outsourcing share. Intuitively, the larger that covariance, the costlier is outsourcing for high-TFP, high value added producers, and therefore the lower the outsourcing gradient. Finally, contractors' productivity wedge is set to match the outsourcing wage penalty.

We define the loss function as

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{N}\sqrt{\sum_i \left(\frac{h_i(\boldsymbol{\theta}) - \hat{h}_i}{\hat{h}_i}\right)^2},$$

where $\boldsymbol{\theta}$ is the vector of parameters to be calibrated, $\hat{h}$ are the targetted moments in the data, and $h(\boldsymbol{\theta})$ are the corresponding moments in the model. We compute the model as described in Section G.3 except that $\{\overline{M}^c, M^g\}$ are treated as given. Instead, we iterate in the outer loop over μ_z and τ to ensure the model matches the average and outsourcing wage penalty. The loss function is defined over the variance-covariance matrix $\mathbf{\Sigma}$ and its associated three moments.

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 $\boldsymbol{\theta}^{j+1} = \boldsymbol{\theta} - \gamma_j \nabla \mathcal{L}(\boldsymbol{\theta})$, where the endogenous step

[36] This normalization has no consequence as the model is scale-free.

size follows the Barzilai–Borwein method. Namely, for $j > 1$,

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

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. The gradient descent is implemented in Julia and parallelized over 3 CPUs. Our starting point for the variance of TFP and outsourcing cost are the variance of in-house and outsourced TFP respectively. For the covariance between the two, we set it up to zero. The descent is run on a standard laptop and takes about 20 minutes to converge.

H.4 Identification

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

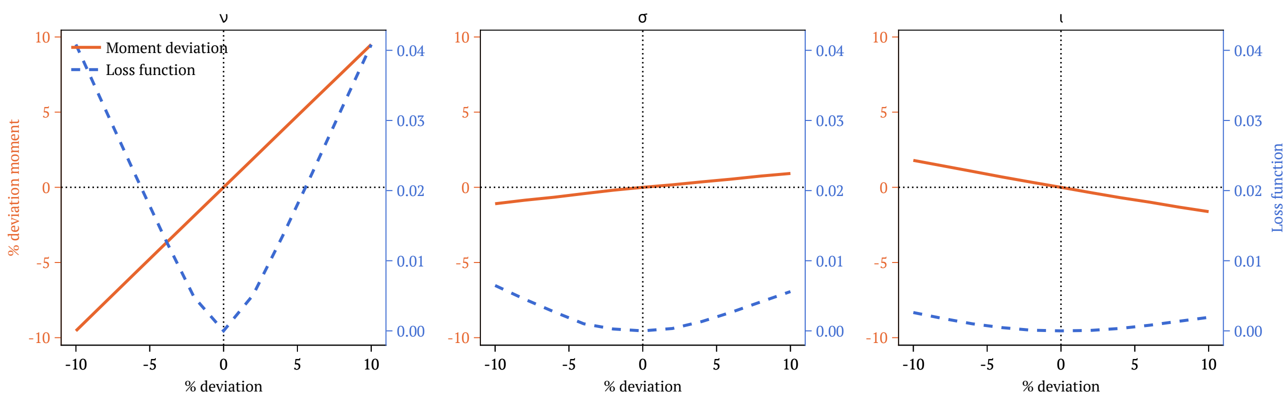

Note: Numerical local identification of parameters. Solid orange line: percentage deviation of targeted moment relative to moment at estimated parameters, as a function of percentage deviation of parameter. Mapping as per Table 1. Dashed blue line: percentage deviation of loss function relative to loss function at estimated parameters, as a function of percentage deviation of parameter.

H.5 Model fit and over-identification

We verify whether the estimated model accounts for selection into outsourcing. We have targeted the cross-sectional OLS coefficient in column (2) from Table 4, Online Appendix F.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 22(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

Figure 22: Selection into outsourcing and the productivity effect in the estimated model.

(a) Selection into outsourcing.

(b) Productivity effect.

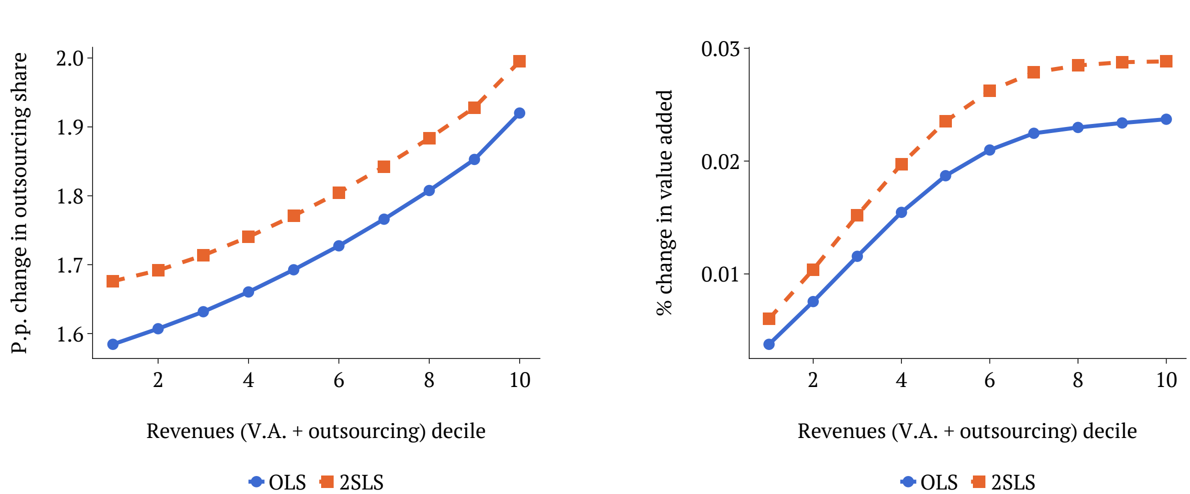

Note: Panel (a): OLS (solid blue) and 2SLS (dashed orange) coefficients in the estimated model. OLS coefficient computed by projecting the change in the outsourcing share on the change in log value added, following a one standard deviation Δz increase in z and the corresponding change in $\Delta \varepsilon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

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 unambiguously 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 to 3.52.

In the model, the within-firm OLS coefficient conflates the change in revenue productivity z with the associated change in outsourcing costs ε. We mimick the instrumental variable strategy from Table 4, 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 share and value added instead of the joint shift in (z, ε). The dashed orange line displays our results. Consistent 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 and 3.55 depending on the incidence of shocks in the model.

Our second over-identification exercise asks whether the estimated model accounts for the untargeted productivity effect. We mimick the OLS and 2SLS coefficients from Table 5 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 22(b) displays the within-firm model counterparts of the OLS and 2SLS coefficients from Table 5, columns (1) and (3), by initial decile of firm value added. In the model, the OLS coefficient ranges from 0.005 to 0.02 depending on the incidence of shocks. The model struggles to generate a high enough 2SLS productivity effect, which remains between 0.01 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.