WORKING PAPER · NO. 2021-05

Outsourcing, Inequality and Aggregate Output

Adrien Bilal and Hugo Lhuillier

JANUARY 2021

Outsourcing, Inequality and Aggregate Output*

Adrien Bilal

Hugo Lhuillier

University of Chicago

Princeton University

December 19, 2020

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 more productive firms either post higher wages along a job ladder to sustain a larger in-house workforce, comprised of many imperfectly substitutable worker types and subject to decreasing returns to scale, or rent labor services from contractors who hire in the same frictional labor markets. Three implications arise: more productive firms are more likely to outsource to save on higher wage premia; outsourcing raises output at the firm level; labor service providers endogenously locate at the bottom of the job ladder, implying that outsourced workers receive lower wages. Using firm-level instruments for outsourcing and revenue productivity, we find empirical support for all three predictions in French administrative data. After structurally estimating the model, we show that the rise in outsourcing in France between 1997 and 2007 increased aggregate output by 1% and reduced the labor share by 3 percentage points.

*Adrien Bilal: adrienbilal@uchicago.edu. Hugo Lhuillier: hugo.lhuillier@princeton.edu. The authors thank Rodrigo Adao, Damien Capelle, John Grisgby, Gordon Hanson, Johnathan Hazell, Erik Hurst, Oleg Itskhoki, Simon Mongey, Stephen Redding, Richard Rogerson, Esteban Rossi-Hansberg, Rob Shimer, Nathan Zorzi, and numerous seminar participants for valuable comments. All errors are ours.

Introduction

Outsourcing is fundamentally changing the nature of the labor market. During the last two decades, firms have been increasingly contracting out a vast array of labor services, such as security guards, food and janitorial services. While workers in these occupations enjoy high wages when working for high-paying traditional employers, they earn much less when employed by a contractor firm. Domestic outsourcing therefore has distributional consequences that adversely impact relative wages of outsourced workers. However, to the extent that firms scale up more efficiently by contracting out certain activities, outsourcing generates aggregate output gains that all workers may share in. Despite the prevalence of outsourcing in the labor market, there is little guidance to trace out its determinants and effects. Why do firms outsource? How can low-paying contractor firms co-exist with high-paying traditional employers? How does outsourcing change aggregate production and its split between workers and firms?

In this paper, we propose answers to these questions in three parts. First, we build an analytic theory of domestic outsourcing based on imperfect rent-sharing in the labor market. Second, using administrative data from France, we offer new empirical evidence that confirms the distributional and productivity effects of outsourcing that our theory highlights. Third, we structurally estimate our general equilibrium model and quantify the effects of outsourcing on aggregate output and the labor share.

Specifically, in the first part paper, we contribute a framework to study the emergence of outsourcing. Conceptually, such a framework should have three key properties. First, firms should have well-defined boundaries. Second, not all workers should be equally exposed to outsourcing. Third, seemingly identical workers should earn different wages at different employers.

To set the stage for our analysis of outsourcing, we start with an environment that features these properties, but no outsourcing yet. Firms have heterogeneous productivities, and are subject to decreasing returns to scale in revenue, either due to downward-sloping demand or to technological reasons. Firms hire workers of different skills in segmented labor markets, who enter as imperfect substitutes in the production process and search for employment opportunities on and off the job. Frictions in skill-specific labor markets give rise to rent-sharing between workers and firms. Firms exploit their monopsony power over workers, leading to wage dispersion, and firms with a larger target size post higher wages. Together, these elements shape the incentives to outsource at the firm level. However, they also depart from the assumptions of constant returns to scale and perfect substitution between workers that are traditionally

imposed to retain traction in wage-posting models.

We show that it is in fact enough to place much less structure on the revenue function to achieve tractability. Our sufficient condition is that the revenue function be supermodular (exhibits positive cross-derivatives) in firm productivity and employment of each worker skill. It ensures that more productive firms always prefer to hire more and, through standard monotone comparative statics results, guarantees that equilibrium wages rise with firm productivity. Our supermodularity assumption allows for standard revenue functions such as a constant elasticity of substitution (CES) function between worker types encapsulated in a decreasing returns upper tier, either because of downward-sloping residual demand or technology.

We then introduce contractor firms in our environment. Contractor firms hire labor in the same frictional labor markets as traditional firms. Contractor firms then sell labor services of their employees in a competitive labor service market. Traditional firms may buy these labor services at an equilibrium price, or hire workers in-house directly in the frictional labor markets. The price of outsourced labor services reflects both wages paid to workers at contractor firms, as well as a markup which we call the cost of outsourcing. This cost may reflect either the cost of capital or a trade cost encapsulating communication and coordination costs. Finally, we capture the idea that contractor firms' core business is precisely to recruit workers and sell their labor services, while recruiting activities divert resources from the core business of traditional firms. Hence, contractor firms have a comparative advantage in recruiting technology relative to traditional firms whose vacancy cost function exhibits more convexity.

Three main implications emerge. First, traditional firms select into outsourcing. More productive traditional firms have the strongest incentives to outsource and save on labor costs because they pay higher wages to attract and retain a larger workforce. Less productive traditional firms who pay lower wages prefer to hire in-house and avoid the outsourcing markup. Second, outsourcing allows traditional firms to sidestep labor market frictions and scale up: employment and revenues increase discretely when traditional firms outsource. Thus, outsourcing leads to a positive total factor productivity effect at the firm level. Third, large wage gaps emerge between traditional and contractor firms. Contractor firms can afford to pay the lowest wages in the economy because they have a comparative advantage in hiring technology. As a result, when a traditional firm outsources its workforce and its former workers transition to a contractor firm, their wages drop from the top of the job ladder to the bottom. The resulting wage losses for outsourced workers capture the distributional effect.

In the second part of the paper, we test the main implications of our theory using admin-

istrative data from France for 1997 to 2007. 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 expenditures on outsourcing at the firm level as expenditures on workers who are not employees of the firm, but are at least partially under the legal authority of the purchasing firm.¹ In the aggregate, expenditures on outsourcing represented 6% of the aggregate wage bill in 1997, before rising to almost 11% in 2007.

First, we test selection into outsourcing. To do so, we investigate the relationship between firm-level value added and the outsourcing share, defined as outsourcing expenditures out of all labor costs including outsourcing. To isolate the effect of revenue productivity from confounding factors such as Information Technologies (IT) improvements that may affect a particular firm’s propensity to outsource, we construct an instrument at the firm level. We interact initial firm-level export shares with changes in foreign demand across 4-digit industries and countries (Hummels et al., 2014). We find that a 1 standard deviation increase in value added leads to a 0.2 to 0.4 standard deviation increase in the outsourcing share. We conclude that firms indeed select into outsourcing.

Second, we test whether a decline in the cost of outsourcing leads firms to scale up. To do so, we investigate the reverse relationship, running from the outsourcing share to firm-level value added. To isolate a decline in firm-level outsourcing costs from revenue productivity, we need a second instrument for firms’ outsourcing share. We interact initial firm-level occupation shares with changes in aggregate spending on outsourcing at the occupation level. We find that a 1 standard deviation increase in the outsourcing share leads to a 0.1 to 0.2 standard deviation increase in value added. We conclude that outsourcing indeed has a positive total factor productivity effect at the firm level.

Third, we confirm the wage penalty from outsourcing in France with an event study design. We define outsourcing events building on Goldschmidt and Schmieder (2017), based on changes in occupational shares, increases in outsourcing expenditures and joint mobility of clusters of workers. We find that job switchers in an outsourcing event lose 12% of their pre-event wage relative to workers at the firm who are not in the outsourcing event but also switch employers. We conclude that outsourcing indeed has distributional consequences on relative wages of workers.

¹We must stop our analysis in 2007 due to a substantial change in the administrative data, as well as in the questions asked in the survey that we use to measure outsourcing. After 2008, it is unfortunately not possible to measure expenditures on outsourcing, nor to reliably identify contractor firms based on their industry codes.

In the third part of the paper, we develop and structurally estimate a quantitative version of the framework with two skill types, before investigating the role of the rise in outsourcing for inequality and output. The main additions are a flexible curvature in traditional firms' vacancy cost function, and firm-level outsourcing costs leading to mixing in the outsourcing decision at all scales, while preserving selection into outsourcing on average.

We estimate the model with a Method of Simulated Moments (MSM) estimator. We discipline selection into outsourcing using moments that connect to our reduced-form estimates of the productivity effect. Importantly, we do not target dispersion in firm wage premia—informed solely through worker flows and productivity dispersion—nor the wage penalty from outsourcing. We show numerically that the model is locally identified.

The third part of this paper quantifies the race between the productivity and the distributional effects in the aggregate. We start with a set of validation exercises using non-targeted moments that inform the redistributive consequences of outsourcing. First, the standard deviation of firm wage premia predicted by the model is 0.15, against 0.14 in the data. Second, the predicted outsourcing wage penalty is 17% in the estimated model, and we cannot reject that it is equal to its empirical counterpart at the 5% level. Third, when using occupation and industry codes to measure the economy's employment share in contractor firms, we obtain 3.8% in the data in 1997, against 4% in the model. Together, these observations support the estimated model's ability to account for the distributional implications of outsourcing.

We then conduct our main counterfactual exercise. We change the cost of outsourcing such that outsourcing expenditures mimic the rise seen in France between 1997 and 2007. Aggregate output rises by 1%, as labor is effectively reallocated to the most productive firms in the economy. However, these productive gains are unevenly distributed. Low-skill workers, who are particularly exposed to outsourcing, are increasingly employed at contractor firms who pay low wages. In addition, their wages decline even at traditional employers. This general equilibrium response arises because traditional employers now face weaker labor market competition for workers. Traditional employers can easily poach workers from contractor firms, while shielded from wage competition from firms previously at the top of the job ladder, who now outsource and left the labor market. Put together, these results imply that the labor share declines by 3 percentage points, and aggregate labor income drops by 2%. We conclude that outsourcing leads to some, though modest, positive productivity effects, and that these gains benefit firm owners and deteriorate workers' labor market prospects.

Literature. This paper relates to several strands of literature. The first is the empirical literature that studies the wage and employment effects of outsourcing. Goldschmidt and Schmieder (2017) and Drenik et al. (2020) document that outsourced workers experience large wage declines in Germany and Argentina, respectively. Katz and Krueger (2017) document a rise in alternative work arrangements in the U.S. Bergeaud et al. (2020) highlight that internet broadband expansion lead 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. Bertrand et al. (2020) show that an increase in the supply of contract labor helped Indian firms scale up. We complement this literature by providing a micro-founded theory of outsourcing, and testing its firm-level implications using direct measures of outsourcing expenditures.

Second, our paper connects to the large literature studying how labor market frictions give rise to wage premia across employers. We build on the wage-posting tradition, starting with Burdett and Mortensen (1998), and enriched with multiple worker types by Engbom and Moser (2018). We contribute to this literature by providing sufficient conditions to depart from constant returns to scale and perfect substitutability between workers in production.²

Finally, our paper relates more broadly to the literature on trade in intermediate inputs and international offshoring—see for instance Feenstra and Hanson (1999) and Grossman and Rossi-Hansberg (2008).³ When firms trade intermediate inputs, they contract on a physical good. Under domestic outsourcing, firms contract on a worker’s flow of services, 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 understand how to break the law of one price in the domestic labor market.

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

²See Gouin-Bonenfant (2018) for a wage-posting model with productivity fluctuations. Bargaining models such as Cahuc et al. (2006) also deliver wage premia across employers, though their predictions are less transparent. Helpman et al. (2010) develop a framework with firm wage premia based on worker selection, which however cannot rationalize the findings of Goldschmidt and Schmieder (2017). Finally, models based on compensating differentials such as Card et al. (2018) imply that individuals working at high-paying firms attain exactly the same expected utility as individuals working at low-paying firms, making any distributional effects hard to interpret.

³See also Acemoglu et al. (2015) and Antràs et al. (2017).

1 A theory of wage premia with large firms

We start the exposition of our environment with an economy in which wage premia across employers arise endogenously as the result of the monopsony power employers exert over workers. We add outsourcing to this baseline economy in Section 2, and add quantitative extensions to the framework in Section 4.

1.1 Setup

Time. Time is continuous, and we focus on a steady-state equilibrium.

Workers. There is a unit mass of workers. Each worker is characterized by its exogenous and permanent skill type $s \geq 0$. We assume that types are distributed in the population according to the measure $m_s ds$ with respect to a base measure denoted by ds . This notation allows us to capture both continuous and discrete type distributions without loss of generality.

Workers have linear preferences in income at every point in time, 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 real skill-specific unemployment benefits b_s .

Firms. There is a mass M of active firms in the economy. They are indexed by productivity z with full support in $[\underline{z}, \bar{z}]$, with $\bar{z} \leq +\infty$. Denote by Γ the cumulative distribution function. Assume for simplicity that it admits a finite density, and that $\underline{z}$ is large enough relative to $\sup_s b_s$ so that all matches are viable. A firm with productivity z that hires a mass n_s of workers of skill $s \geq 0$ generates revenue $R(z, \mathbf{n})$, where $\mathbf{n} = \{n_s\}_s$ denotes the vector of employment across worker types. Assume that R is twice continuously differentiable and increasing in each argument.

Labor markets. Labor markets are segmented by skill s . A labor market consists of an equilibrium distribution of skill-specific wage offers, and job searchers. Unemployed workers of skill s sample wage offers randomly at Poisson intensity λ_s^U . Employed workers of skill s sample wage offers randomly at Poisson intensity λ_s^E from the same distribution. Employed workers can break their current contract to accept a new wage offer. Existing matches are destroyed at Poisson rate δ_s . Thus, a match ends either when it is exogenously destroyed, either when the worker accepts a wage offer.

Wage contracts. To attract and retain workers, firms optimally post wage offers in every skill-specific labor market. Every firm is endowed with a unit measure of vacancies for every skill s to which they attach the same skill-specific wage offer. Firms fully commit to a single wage per skill once posted.

1.2 Labor market transitions and labor supply curve

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

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

The value of being employed at wage w , $V_s(w)$, is increasing with the wage w , so that workers behave as income maximizers: they always accept higher wage offers while employed. Equating the value of being employed to the value of being unemployed defines the reservation wage $\underline{w}_s$, given in Appendix A.1.

The equilibrium distribution of wages of employed workers, denoted $G_s(w)$, determines the labor supply curve of each firm. Because workers optimally switch to higher-paying jobs when they sample one, we can relate the wage offer distribution $F_s(w)$ to the wage distribution of employed workers $G_s(w)$ using worker flows. We show in Appendix A.2 that the distribution of wages of employed workers writes

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

From equation (1) we may characterize the number $N_s(w)$ of employed workers per wage offer in the interval $(w - \varepsilon, w]$ for skill s .⁴ We obtain

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

⁴ $N_s(w)$ is equal to the limit of the ratio $\frac{G_s(w) - G_s(w - \varepsilon)}{F_s(w) - F_s(w - \varepsilon)}$ when $\varepsilon \rightarrow 0$, times the number of employed workers $m_s - u_s$.

where $e_s = \frac{\lambda_s^U m_s}{\delta_s + \lambda_s^U}$ is the mass of employed workers of skill s .

Crucially, $N_s(w)$ is non-decreasing in the wage w . $N_s(w)$ thus defines the upward-sloping labor supply curve faced by firms. It depends on the equilibrium distribution of wage offers in the economy, $F_s(w)$. We now turn to firms' decision problem to characterize this distribution.

1.3 Wage and employment distributions

Since firms post a unit measure of vacancies, the number of workers per firm posting w , $n_s(w)$, is simply related to the number of workers employed at every wage by

$$n_s(w) = \frac{N_s(w)}{M}.$$

Then, when the discount rate is low enough, firms choose their wage offers $\{w_s(z)\}_s$ to maximize their flow profits.⁵ Firms take as given how their size depends on their wage offer through their upward-sloping labor supply curve given in equation (2). Flow profits are given by:

$$\begin{aligned} \pi(z) &= \max_{\{w_s\}_s} R\left(z, \{n_s(w_s)\}_s\right) - \int w_s n_s(w_s) ds \\ \text{s.t. } n_s(w) &= \frac{(1 + k_s)e_s}{M(1 + k_s(1 - F_s(w)))(1 + k_s(1 - F_s(w^-)))} \quad \forall s. \end{aligned} \quad (3)$$

Unless the distribution $F_s(w)$ can be characterized more precisely, the problem in equations (3) is intractable. Existing wage-posting models such as Burdett and Mortensen (1998) and Engbom and Moser (2018) have leveraged two key simplifying assumptions to make progress. Under constant returns and perfect substitutability of workers in production, $R(z, \mathbf{n}) = z \int n_s ds$, the problem (3) can be split at the match level. Once decoupled across matches, it is straightforward to see that (3) exhibits a single-crossing property. This structure implies that wages are increasing in productivity z , which in turn allows to solve for the distribution of wage offers in terms of the equilibrium wage policy and the exogenous productivity distribution, $F_s(w_s(z)) = \Gamma(z)$.

Studying outsourcing requires however a well-defined boundary of the firm as well as possible interactions between workers in production. Thus, we relax both assumptions of constant returns and perfect substitutability. We now state our main sufficient condition on the revenue function R under which we may still rank wages by firm productivity using tools from the monotone comparative statics literature.

⁵We fully derive the formulation in equation (3) from the dynamic problem of the firm in Appendix A.3.

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

Given that R is twice continuously differentiable, Assumption (A) is equivalent to imposing 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) guarantees that as productivity rises, firms always prefer to hire more labor of every type, through two channels. First, the direct effect of a productivity increase incentivizes more hiring of that labor type because productivity and labor of any type are complements in levels. Second, Assumption (A) also ensures that the firm never prefers to lower employment of a given labor type as it hires more labor of another type, holding productivity fixed.

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 proportional deviations. The following examples show that our supermodularity assumption allows for workers to be complements or substitutes in production in the usual sense.

Example 1. Workers have a CES utility function over the M differentiated varieties i with elasticity of substitution $\sigma > 1$, leading to a downward-sloping residual demand for every variety $p_i = Aq_i^{-1/\sigma}$, where A is a general equilibrium constant. Each variety i is produced by a single firm. Firms produce with a CES production function with elasticity of substitution $\eta \in [0, +\infty]$, $y(z, \mathbf{n}) = z \left(\int (a_s n_s)^{1-\frac{1}{\eta}} ds \right)^{\frac{\eta}{\eta-1}}$. The revenue function is $R(z, \mathbf{n}) = z^{1-\frac{1}{\sigma}} \left(\int (a_s n_s)^{1-\frac{1}{\eta}} ds \right)^{\frac{\eta(\sigma-1)}{\sigma(\eta-1)}}$.

Supermodularity then requires $\sigma > \eta$. 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.

Example 2. Suppose that instead of varieties, there is a homogeneous final good produced by all firms and used as the numeraire but there are technological decreasing returns to scale in production: $y(z, \mathbf{n}) = z \left(\int (a_s n_s)^{1-\frac{1}{\eta}} ds \right)^{\frac{\rho\eta}{\eta-1}}$. The revenue function is then $R(z, \mathbf{n}) = y(z, \mathbf{n})$, and supermodularity requires $\frac{1}{1-\rho} > \eta$. With typical estimates of ρ close to 0.8, supermodularity is again compatible with standard parametrizations.

We conclude that not only does Assumption (A) ensure that more productive firms have

incentives to be larger, but it is also compatible with a relatively general class of revenue functions. We impose Assumption (A) henceforth in the paper, and now formalize how it ensures that more productive firms post higher wages.

Proposition 1. (*Wage ranking*)

In equilibrium, wages $w_s(z)$ are strictly increasing with firm productivity z . The wage function is jointly continuous, and the wage offer distribution satisfies

$$F_s(w_s(z)) = \Gamma(z).$$

With Proposition 1 at hand, the distribution of workers across skills and firms is fully determined by workers' transition rates in their labor market, and the productivity ranking of firms. Proposition 2 summarizes our result.

Proposition 2. (*Employment distribution*)

The number of workers of skill s hired by firm z is given by

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

Notice that firm size in Proposition 2 depends only on the ranking of firms, $\Gamma(z)$, because firm size is fully determined by worker flows up the job ladder. This stark result arises because we do not let firms choose how many vacancies to post—which is equivalent to a vacancy cost with infinite curvature. We introduce endogenous vacancies with a cost function with finite curvature in our quantitative extensions in Section 4, so that firm size also reflects the marginal product of labor.

To understand why $k_s = \lambda_s^E/\delta_s$ enters in Proposition 2, recall that workers start from unemployment and initially accept relatively low-paying jobs. While searching on the job, they accept any offer above their current wage. The speed at which they climb the job ladder in their market depends on the frequency at which they receive wage offers, the job-finding rate λ_s^E . They fall down the job ladder back into unemployment at rate δ_s . On net, the allocation of workers along the job ladder depends on the ratio $k_s = \lambda_s^E/\delta_s$.

Building on Propositions 1 and 2, we characterize the solution to firms' wage-posting problem (3) and solve explicitly for the wage distribution.

Proposition 3. (*Wage distribution*)

Equilibrium wages are given by

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

where n is given in Proposition 2.

The expression in Proposition 3 captures the logic of the job ladder. Productive firms raise their wages to poach workers from lower-productivity firms in order to attain their target size. The equilibrium value of a worker to these lower-productive 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, and integrating up to productivity z .

Having characterized the emergence of wage premia across firms in our baseline economy, we are now ready to introduce outsourcing.

2 A theory of outsourcing

In this section, we enrich our basic environment with contractor firms that provide outsourcing services and characterize how they affect the economy. The rest of the environment remains identical to Section 1.

2.1 Contractor firms

There is a continuum of identical contractor firms in every skill market s . To make the distinction clear, we now call firms that produce a consumption good “traditional firms”. Contractor firms hire workers in the same frictional labor market as final good producers, also by posting wages. A given contractor firm hires in a single skill market s , and faces constant returns in production. There is perfect competition in all rental markets for labor services. The equilibrium rental price of one efficiency unit of labor of skill s is denoted p_s .

We model a wedge $\tau_s < 1$ between the price p_s and the revenue that contractor firms earns when a traditional firm buys one unit of labor from them. Contractor firms earn profits

$$\pi^C(w) = (\tau_s p_s - w)n_s(w).$$

We propose two micro-foundations for τ_s , detailed in Appendix A.7. Regardless of the micro-foundation, the wedge τ_s is an exogenous parameter that captures how costly it is to outsource workers. In our first micro-foundation, τ_s simply reflects a parameter, the inverse of an iceberg trade cost between contractor firms and traditional firms. This trade cost captures the idea that communication, monitoring and coordination between the traditional firm and outsourced workers may be more difficult when workers are employees of another firm. As a result, some efficiency units of labor are lost. In our second micro-foundation, contractor firms combine a small amount of capital and labor according to a Cobb-Douglas production function. τ_s then simply encapsulates the equilibrium price of capital.

Finally, we assume that contractor firms operate a different hiring technology than traditional firms. Our goal is to capture the idea that the core activity of contractor firms is precisely to hire workers and sell their labor services to other firms. In contrast, the core activity of traditional firms is to produce a particular good, therefrom hiring activities divert productive resources. Specifically, we assume that contractor firms face a free entry condition:

$$0 \geq \max_w \left(\tau_s p_s - w \right) n_s(w). \quad (4)$$

Given constant returns for contractor firms, free entry is equivalent to costless recruitment activities.

2.2 Traditional firms and outsourcing

Traditional firms now face an additional possibility to buy labor. They may still hire workers in-house in a frictional labor market. The other possibility is now to rent labor services. Formally, their decision problem becomes

$$\pi(z) = \max_{\{n_s\}_s, \{o_s\}_s \in \{0,1\}^S} R(z, \{n_s\}_s) - \int \left[(1 - o_s) w_s(n_s) + o_s p_s \right] n_s ds, \quad (5)$$

where we have used the inverse function of $n_s(w)$, $w_s(n)$, to best contrast the decision to hire in-house with the decision to outsource. The indicators $o_s \in \{0, 1\}$ indicate whether a traditional firm outsources skill s . With the notation in equation (5), n_s denotes in-house labor if $o_s = 0$, and denotes outsourced labor if $o_s = 1$.

If the traditional firm hires in-house ($o_s = 0$), it faces an upward-sloping labor supply curve embedded in the increasing function $w_s(n)$. Thus, a traditional firm with a large target size n_s

ends up paying high wages in-house. In contrast, if the traditional firm outsources ($o_s = 1$), it faces a flat labor supply curve at price p_s . However, since the price p_s reflects both the wage paid to employees of contractor firms as well as the wedge τ_s , the price of outsourcing may exceed in-house wages when traditional firms target a small size n_s . When the traditional firm targets a large size n_s however, outsourcing may become more advantageous due to the upward-sloping wage premia curve $w_s(n)$. We formalize this discussion in the following Proposition.

Proposition 4. (*Outsourcing*)

Consider a skill s . Suppose that τ_s is high enough that there is some outsourcing in equilibrium. Then:

1. Contractor firms pay the reservation wage $\underline{w}_s$;
2. The price of outsourcing is $p_s = \frac{w_s}{\tau_s}$;
3. There exists a threshold productivity $\hat{z}_s$, such that outsourcing occurs if, and only if, $z \geq \hat{z}_s$;
4. The highest wage in the economy is capped by the price of outsourcing: $w_s(\hat{z}_s) \leq \frac{w_s}{\tau_s}$;
5. If, in addition, $\frac{\partial^2 R}{\partial n_s^2} < 0$, then the previous inequality is strict, and revenue and labor inputs discretely jump up as firms outsource: $R(\hat{z}_s^+, \mathbf{n}^*(\hat{z}_s^+)) > R(\hat{z}_s^-, \mathbf{n}^*(\hat{z}_s^-))$ and $n_s(\hat{z}_s^+) > n_s(\hat{z}_s^-)$.

Proposition 4 characterizes how outsourcing shapes the labor market. First, contractor firms pay the lowest wage in the economy, the reservation wage $\underline{w}_s$. This result follows immediately from the zero-profit condition (4) for contractor firms. If any wage $w > \underline{w}_s$ was posted in equilibrium by a contractor firm, it could make positive profits by lowering their wage offer by a small amount, thereby contradicting the zero-profit condition. This result reveals that outsourcing has distributional consequences by reallocating workers towards the lowest-paying firms in the labor market.

Second, the price of outsourcing is a simple markup $1/\tau_s > 1$ over the reservation wage $\underline{w}_s$. This result again follows directly from the free-entry condition of contractor firms (4).

Third, we obtain selection into outsourcing. Since more productive firms have a larger target size, they pay higher wage premia if they hire in-house. Thus, highly productive firms have the strongest incentives to outsource.

Fourth, the price of outsourcing is an effective wage cap in the labor market. Firms never pay wages above the price of outsourcing because they always have the option to outsource.

Only the most productive firms outsource. Thus, outsourcing removes the highest-paying jobs from the job ladder. This result highlights that not only does outsourcing reallocate workers towards low-paying jobs, outsourcing also removes workers' best options from the labor market, thereby reinforcing its distributional consequences.

Fifth, the amount of efficiency units of labor hired by a traditional firm jumps up upon outsourcing. This increase in effective firm size arises because traditional firms equate the marginal cost of labor to its marginal product. When outsourcing, traditional firms switch from a convex to a linear labor cost curve. Thus, they are able to scale up, leading to a productivity effect.

Proposition 4 does not only characterize the distributional and productivity effects of outsourcing, it can be readily used to interpret the outsourcing wage penalty in the data. Consider a worker of skill s at traditional firm z that is close to the outsourcing cutoff, $z < \hat{z}_s < z + dz$ for a small dz . Suppose the traditional firm experiences a small idiosyncratic increase in productivity dz' such that $z + dz' > \hat{z}_s$. Then, the traditional firm lays off their workers of skill s and hires labor services in the rental market. Suppose either that some of these workers are immediately re-hired by contractor firms as in Goldschmidt and Schmieder (2017). Outsourced workers then experience a discrete wage drop from $w_s(\hat{z}_s)$ to $\underline{w}_s$.

2.3 Equilibrium

To complete the description of the equilibrium, consider the wage offer distribution. Proposition 4 states that every contractor firm posts the reservation wage. With some outsourcing in equilibrium, the wage offer distribution thus starts with a mass point at the reservation wage. Traditional firms $z \in [\underline{z}, \min_s \hat{z}_s]$ then behave similarly to the no-outsourcing economy. For these firms, $o_s = 1$, and so the only change relative to Propositions 1, 2 and 3 stems from the number of workers they can attract and retain with a wage offer. These traditional firms are able to poach workers from contractor firms at the bottom of the job ladder, and no longer face competition from the most productive traditional firms who now outsource their labor services. Thus, their equilibrium rank in the job ladder is given by

$$\Upsilon_s(z) = \frac{\mathcal{V}_s + M\Gamma(z)}{\mathcal{V}_s + M\Gamma(\hat{z}_s)}, \quad \forall z \leq \hat{z}_s, \quad (6)$$

where $\mathcal{V}_s$ denotes the measure of wage offers by contractor firms. Their equilibrium size and wage then follows from replacing $\Gamma(z)$ in Propositions 2 and 3 by $\Upsilon_s(z)$.

The market clearing condition for labor services determines how many contractor firms find it profitable to operate:

$$M \int_{\hat{z}_s}^{\bar{z}} n_s(z) d\Gamma(z) = \iota_s n_s^{\text{Out}} \frac{\mathcal{V}_s}{\mathcal{V}_s + M\Gamma(\hat{z}_s)}, \quad \forall z \leq \hat{z}_s, \quad (7)$$

where, when evaluated at $z > \hat{z}_s$, $n_s(z)$ is the demand for outsourced labor services which follows from the traditional firm's first-order condition. $n_s^{\text{Out}} = \frac{e_s}{1+k_s}$ indicates how many workers contractor firms manage to recruit and retain per wage offer, from equation (2). ι_s is a dummy-type variable that depends on the type of micro-foundation we choose for outsourcing. ι_s is equal to τ_s when we micro-found the cost of outsourcing with an iceberg trade cost. ι_s is equal to 1 when we micro-found the cost of outsourcing with capital use.

2.4 Testable implications

The theory of outsourcing we have laid out delivers several testable implications that follow from Proposition 4. We discuss them below.

Firms with larger revenue spend relatively more on outsourcing services. Selection into outsourcing (Proposition 4.3) implies that more productive firms are more likely to outsource. This observation implies that firms with a larger revenue spend a larger share of their overall labor costs—inclusive of outsourcing expenses—on outsourcing.

The positive relationship between revenue and the outsourcing share lies at the core of our theory's structure. Yet, it only indirectly informs the productivity effect, and does not characterize the distributional effect. Thus, we also propose more direct tests of the productivity and distributional effects.

Firms with lower costs of outsourcing produce more. To understand how the productivity effect emerges in our theory, consider a firm z just below an outsourcing threshold $\hat{z}_s$. Anticipating on Section 4 in which we introduce idiosyncratic outsourcing costs, suppose that the idiosyncratic cost of outsourcing faced by firm z drops by a small amount. Then firm z will switch into outsourcing skill s . As per Proposition 4.5, its share of expenditures on outsourcing increase together with its revenue. Together, these observations imply that when firms are subject to outsourcing cost shocks leading to a rise in outsourcing expenditures, their revenues rise too.

Wage penalty from outsourcing. To test the distributional effect, we focus on the wage penalty from outsourcing, measured as the wage loss when workers transition from a traditional firm to a contractor firm. Proposition 4 implies that contractor firms pay the lowest wages in the labor market. Again anticipating Section 4, when traditional firms experience an idiosyncratic outsourcing cost shock that leads them to outsource and some of their workers transition into contractor firms, the wages of these workers drop.

Having described our theory of outsourcing and its implications, we now implement our tests thereof.

3 Reduced-form evidence

In this section, we first describe our data. Second, we discuss aggregate trends in outsourcing in France. Then, we proceed to test our three main predictions: selection into outsourcing, the productivity effect, and the distributional effect. Additional details are in Appendix C.

3.1 Data

We use a combination of administrative and survey data for France between 1996 and 2007. Our first data source are annual tax records for the near-universe of French firms (FICUS) that report balance sheet information. Among others, we observe total employment and wage bill, sales and purchases of intermediate inputs, from which we construct value added. However, this balance-sheet information does not break down purchases of intermediate inputs finely enough to isolate outsourcing expenditures.

Our second data source is a large, mandatory annual firm-level survey that provides a more detailed breakdown of firms purchases of intermediate inputs (EAE). Surveyed firms report expenditures on “external workers”. External workers are employees of another firm, but that fall under a contracting agreement with the surveyed firm. Importantly, these workers 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. Finally, firms remain in the survey once they enter, which allows us to leverage the panel dimension of this data.

Our third data source consists of employer tax records that provide information on labor market outcomes for French workers (DADS). We use repeated cross-sections that cover the universe of French workers to construct employment and wages at the firm-occupation-year

level. We also use a 4% representative panel that tracks workers throughout their labor market histories to study the wage penalty of outsourcing.

Our fourth data source are customs records for the universe of trade transactions. We observe imports and exports crossing the French border, at the product-country-firm-year level. We use this data to construct export demand shocks at the firm-level and generate variation in firm scale.

We link these four data sources together using a common firm identifier. For our main empirical exercises at the firm level, we aggregate years into three periods 1996-1999, 2000-2002, 2003-2007 and keep only firms with at least ten in-house employees in order to limit measurement error in outsourcing expenditures. We stop our main analysis in 2007 because of a large change in classification, including a change in the structure of the firm-level survey that prevents us from measuring outsourcing expenditures in subsequent years. Our final sample consists of 173,547 firm-periods.

3.2 Aggregate trends in outsourcing

We start by asking how much did outsourcing rise in France. Figure 1(a) shows that outsourcing expenditures as a fraction of the aggregate wage bill almost doubled in the decade that we study: it increased from 6% in 1997 to 11% in 2007. While we cannot reliably measure outsourcing expenditures in the following years, the trend suggests that the rise in outsourcing may have continued in subsequent years.

Does this rise in aggregate outsourcing expenditures translate into an aggregate increase in the employment share of contractor firms? Answering this question requires to have a reliable way to identify contractor firms in the data. It also requires to separate employees at contractor firms whose labor services are rented out to other firms from other employees such as managers. To make progress, we follow Goldschmidt and Schmieder (2017) and rely on industry and occupation codes to detect contractor firms and service workers at contractor firms. To define a low-skill contractor firm, we use industry codes that specifically label firms as providing food, security, cleaning or general administrative services to other firms. To define a high-skill contractor firm, we use industry codes that label firms as providing accounting, law or consulting services to other firms. An important caveat to that approach is that it may miss any firm that is a contractor firm, but does not fall into those specific industry codes, while our measure using outsourcing expenditures is not subject to this limitation.

Figure 1: Outsourcing in France.

(a) Outsourcing expenditures relative to wage bill.

(b) Fraction of employment at contractors.

(c) Fraction of service employment at contractors.

Figure 1(b) shows the fraction of workers who are employees at low-skill contractor and high-skill contractor firms. The employment share at low-skill contractor firms rises from 5% in 1997 to 9% in 2007. At high-skill contractor firms, the employment share increases more modestly from 2% to 2.5%. The rise in employment shares at contractor firms closely tracks the outsourcing expenditure share.

Of course, the services provided by a food contractor firms are likely to require mostly workers in occupations related to food services. To assess whether employment in such occupations has become increasingly concentrated in food contractor firms, Figure 1(c) shows the share of employment in service occupations—defined as food, security, cleaning or general administrative occupations for low-skill employment, and accounting, law or consulting occupations for high-skill employment—at contractor firms. It rises steadily over time, indicating that workers in service occupations have reallocated towards contracting firms over time. Thus, the rise in outsourcing arises not only in expenditures, but also in employment.

3.3 Selection into outsourcing

We start by testing the core prediction of our theory: selection into outsourcing. Following our discussion in Section 2.4, firms with larger revenues should spend relatively more on outsourcing services. We use value added as our main measure of firm revenues because our theory abstracts from the use of other intermediate inputs in production. We define the outsourcing share of a firm as its expenditures on external workers divided by the sum of its expenditures on labor, defined as the sum of its wage bill plus expenditures on external workers.

Figure 2(a) plots the outsourcing share by decile of value added. The solid blue line shows raw data, while the dashed orange line displays the same relationship after removing 3-digit industry and time period fixed effects from both variables. In both cases, Figure 2(a) reveals that high value-added firms tend to outsource more. A firm in the first decile of value added spends less than 4% of its labor costs on outsourced labor, while a firm in the tenth decile of value added spends over 8%. The residualized relationship reveals an S-shaped pattern, steeper at intermediate levels of value added and flatter at the extremes.

While our theory has no unambiguous prediction as to whether the positive relationship between the outsourcing share and value added should also hold when using in-house employment, Figure 2(b) reveals that it is also the case empirically. Firms with more in-house employees also spend relatively more on outsourced labor.

Of course, the relationship depicted in Figure 2 could be the result of unobserved firm-level heterogeneity that would both affect productivity and the ability to outsource. IT-intensity would be an obvious example. Therefore, we turn to a regression design in order to assess the robustness of our results. We consider econometric specifications of the following form:

$$S_{ft} = \alpha_t + \beta_f + \gamma \log VA_{ft} + \varepsilon_{ft} \quad (8)$$

Figure 2: Outsourcing share by value added and size deciles.

where f indexes firms and t indexes time periods. $S_{ft} = \frac{\mathcal{E}_{ft}}{\mathcal{E}_{ft} + \mathcal{W}_{ft}}$ denotes the outsourcing share, where $\mathcal{E}_{ft}$ is expenditures on external workers, and $\mathcal{W}_{ft}$ the wage bill. α_t is a time period fixed effect, β_f a firm fixed effect, and ε_{ft} a mean zero residual.

Conditioning on firm fixed effects removes any confounding variation that may affect both value added and the outsourcing share, such as average firm-level IT-intensity. Our coefficient of interest, γ , is then identified off how changes in the outsourcing share co-move with changes in value added. However, a further concern may be that firm-level IT-intensity is not constant over time. In that case, changes in firm-level IT-intensity may drive both an increase in value added and the outsourcing share, confounding our interpretation of a positive estimate of γ .

To address this concern and generate quasi-experimental variation in firm value added, we leverage the granularity of our customs data. We follow Hummels et al. (2014) and 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. Our instrument is thus defined as

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

To the extent that export demand growth at the industry-country level is orthogonal to firms'

Table 1: Relationship between firm-level outsourcing shares and firm-level value added. 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 (P25,P75)	(5) OLS (P25,P75)	(6) 2SLS	(7) 2SLS	(8) 2SLS
Log Value Added	1.457*** (0.027)	1.535*** (0.027)	0.679*** (0.055)	2.131*** (0.312)	1.579** (0.609)	3.348** (1.029)		
Log Size							5.662** (1.733)	
Log Labor Prod.								11.461** (3.939)
<i>Fixed Effects</i>								
Year	✓	✓	✓	✓	✓	✓	✓	✓
Industry		✓						
Firm			✓	✓	✓	✓	✓	✓
Obs.	173547	173540	173547	61549	14561	46152	46152	46152
Stand. coef.	0.189	0.200	0.088	0.277	0.205	0.435	0.736	1.491
1 st -stage F-stat.	.	.	.	.	.	273	193	61

Standard errors in parenthesis, clustered by firm. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Variables winsized at 5% level. Includes period dummies. Instrument: shift-share of export demand growth by 4-digit industry, projected by firm using firm-level export shares in first period. Restriction in columns (4-5): residual value added (net of year and firm fixed effects) between 25th and 75th percentile of its distribution.

idiosyncratic ability to outsource, $Z_{f,t}$ is a valid instrument for firm value added.

Table 1 displays our results. Columns (1-2) show the regression analog of Figure 2(a) for comparison purposes. The estimate in columns (2) implies that a 100 log points increase in value added is associated with a 1.54 percentage point increase in the outsourcing share. Once standardized, this coefficient implies that a one standard deviation increase in value added is associated with a 0.2 standard deviation increase in the outsourcing share. Column (3) then introduces firm fixed effects. Our estimate drops to 0.68. This reduction is consistent with two possible explanations. First, unobserved heterogeneity such as IT-intensity is partly driving the larger values in columns (1-2). Second, changes in value added are relatively more frequent at the extremes of the value added distribution, perhaps due to mean-reversion. Such a feature would imply that the estimate in column (3) is primarily driven by the flat parts of the S-shaped relationship uncovered in Figure 2, leading to a lower estimated value.

To assess whether the difference between our estimates without and with firm fixed effects stems primarily from unobserved heterogeneity or from a nonlinear relationship, column (4) restricts attention to firm-time period pairs such that the change in residual firm-level value

added, after removing firm fixed effects, is between the 25th and 75th percentiles of its distribution. Our point estimate rises to 2.13. This result is consistent with an important role for nonlinearities rather than firm-level unobserved heterogeneity in driving the difference between columns (2) and (3).

Nonetheless, we then introduce our firm-level instrument. Since our instrument only affects exporters, who may not be fully representative of the full population of firms, column (5) first shows our Ordinary Least Square (OLS) estimate on the sample of exporters only. The estimate drops to 1.60, but is not statistically different from the one for the full sample of firms in column (4). Column (6) then shows our Two-Stage Least Square (2SLS) estimate. The estimate rises to 3.35, but is again not statistically different from the one in column (4). In light of the difference between columns (3) and (4), the larger value of our 2SLS is likely to also reflect the local nature of the 2SLS estimate that may concentrate the identifying variation in the steep part of the relationship between value added and the outsourcing share. The first-stage F-statistic is 273, which is well above conventional thresholds for weak instruments. We confirm our results with other metrics of firm performance such as size and value added per worker in columns (7-8).

From Table 1, we thus conclude that firms select into outsourcing, with more productive firms outsourcing more. We now turn to our test of the productivity effect of outsourcing.

3.4 The productivity effect

Having established selection into outsourcing, we turn to our test of the productivity effect. Following our discussion in Section 2.4, firms that outsource more for idiosyncratic reasons should be able to scale up and generate more revenues. Similarly to equation (8) in Section 3.3, we consider econometric specifications of the following form:

$$\log VA_{ft} = \alpha'_t + \beta'_f + \gamma' S_{ft} + \varepsilon'_{ft}. \quad (9)$$

Relative to equation (8), equation (9) interchanges the dependent and independent variables. Of course, comparing the OLS estimates of equations (8) and (9) would simply amount to rescale the conditional correlation between both variables. Relative to equation (8), equation (9) conveys content only if we can isolate variation in the outsourcing share S_{ft} that arises due to exogenous shocks to firms' idiosyncratic cost of outsourcing.

Thus, we construct an instrument for the outsourcing share. Our goal is that it be orthogonal to firm-level changes in revenue productivity z . We leverage differential exposure of firms to

Table 2: Relationship between firm-level outsourcing shares and firm-level value added.

	Log VA			Log VA/N	Log N
	(1)	(2)	(3)	(4)	(5)
	OLS	OLS	2SLS	2SLS	2SLS
	(P25,P75)				
Outsourcing Share	0.007*** (0.001)	0.027*** (0.004)	0.050*** (0.002)	0.026*** (0.001)	0.022*** (0.001)
<i>Fixed Effects</i>					
Firm	✓	✓	✓	✓	✓
Obs.	94336	47168	94336	94336	94336
Stand. coef.	0.031	0.119	0.220	0.114	0.095
1 st -stage F-stat.			16837	16837	16837

Standard errors in parenthesis, clustered by firm. ⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Variables winsorized at 5% level. Includes period dummies. Instrument: shift-share of outsourcing spending growth, projected by firm using firm-level occupation shares in first period. Restriction in column (2): residual outsourcing share (net of firm fixed effects) between 25th and 75th percentile of its distribution. Due to changes in the occupation classification in 2002, we lose the time period 1997-1999 and only have two time periods and thus no time period fixed effects.

service occupations: food, security, cleaning or general administrative occupations. We compute the within-firm employment share for each of these services occupations o in the initial period, ω_{f,o,t_0} . We then interact these firm-level initial employment shares with a measure of the change in aggregate outsourcing spending on occupation o , $\Delta \log \Omega_{o,t,-f}$, net of firm f 's spending.⁶ Our instrument is therefore defined by

$$Z'_{f,t} = \sum_o \omega_{f,o,t_0} \Delta \log \Omega_{o,t,-f}$$

Table 2 displays our results. Columns (1-2) first present OLS estimates for comparability. Column (2) again focuses on non-extreme values of residual changes in the outsourcing share. We find that a 1 percentage point increase in the outsourcing share is associated with a 2.7% increase in value added. This coefficient translates into a 0.12 standard deviation increase in value added for every standard deviation increase in the outsourcing share. As discussed however, there is little additional information in the OLS estimates of Table 2 relative to those in Table 1.

⁶To construct aggregate outsourcing spending on occupation o , we first infer outsourcing expenditures by occupation at the firm level by interacting initial employment shares ω_{f,o,t_0} with firm-level outsourcing expenditures $\mathcal{E}_{f,t}$. We then sum across firms to define $\Omega_{o,t,-f} = \sum_f \omega_{f,o,t_0} \mathcal{E}_{f,t}$. We then difference in time to remove the contribution of time t_0 expenditures and remove firm f 's expenditures to obtain $\Delta \log \Omega_{o,t,-f}$.

Column (3) shows our 2SLS estimate. We find that a 1 percentage point increase in the outsourcing share implies a 5.0% increase in value added. This coefficient translates into a 0.22 standard deviation increase in value added for every standard deviation increase in the outsourcing share. Columns (4-5) reveal that this increase in value added is approximately equally spread between an increase in labor productivity and an increase in in-house employment.

From Table 2, we thus conclude that outsourcing has a positive productivity effect at the firm level. We now turn to our test of the distributional effect of outsourcing.

3.5 The distributional effect

Having established the productivity effect of outsourcing, we turn to our tests of the distributional effect. We propose two exercises. First, we show that contractor firms locate at the bottom of the job ladder by computing the contractor wage premium and related labor market statistics. Second, we use an event study design to measure the wage penalty of outsourcing.

Contractor firms locate at the bottom of the job ladder. According to our theory, contractor firms should pay the lowest wages in the economy, conditional on worker type. Thus, we begin by measuring the wage premium paid by contractor firms, using the industry codes we associate with contractor firms. To do so, we run a two-way fixed effects regression in the spirit of Abowd et al. (1999):

$$\log w_{i,t} = \varphi_i + \psi_{J(i,t)} + \eta_{i,t}. \quad (10)$$

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

As in Engbom and Moser (2018), when λ_s^E/δ_s is independent from s , $b_s = ba_s$, and the revenue function is linear, the wage formula in Proposition 3 is log-additive in a worker effect and a firm effect. In addition, worker mobility is then conditionally random as in Card et al. (2013). Together, these observations imply that equation (10) can be consistently estimated by OLS. Yet, when the aforementioned assumptions are relaxed, our theory's wage equation is not log-additive. In this case, while no longer exact, equation (10) still provides a useful diagnostic device to measure the average wage premium paid by a firm.

Estimating the full distribution of worker and firm effects in equation (10) leads to the well-known limited mobility bias. Hence, we follow the clustering approach developed in Bonhomme

et al. (2019). We group workers and firms each in 50 equally populated groups, based on the unconditional mean worker and mean firm wage. We then estimate equation (10) with OLS at the group level.⁷

We estimate the standard deviation of firm effects to be 0.14.⁸ We compute the mean firm effect for contractor firms that we identify based on industry codes. We find that the mean contractor firm wage premium is -0.12 relative to the mean firm wage premium normalized to 0. Thus, contractor firms pay wages that are almost one standard deviation of firm wage premia below the average firm wage premium. We conclude that, consistently with our theory, contractor firms indeed pay wages towards the bottom of the job ladder.

That contractor firms locate at the bottom of the job ladder also implies that they should hire relatively more workers from unemployment than traditional firms.⁹ We compute the fraction of hires of contractor firms from unemployment using our worker-level panel data.¹⁰ We find that the share of hires from unemployment is 10 percentage points higher for contractor firms relative to other firms. We conclude that, consistently with our theory, contractor firms display hiring patterns that are consistent with being located towards the bottom of the job ladder.

Event study. So far, our analysis of the wage penalty of outsourcing has focused on employers that can be readily identified as contractor firms based on their industry codes. To confirm that our conclusions are not affected by this particular definition, we also conduct an event study to measure the wage penalty of outsourcing that does not rely on industry codes.

We define an “outsourcing event” at firm f and occupation o in year t for worker i if the following conditions are met. First, the employment share of occupation o at firm f drops by at least 25% between year t and year $t + 1$. Second, outsourcing expenditures rise at firm f rise by at least 50% of the corresponding wage bill reduction, between t and $t + 1$. Third, worker i transitions to a firm f' within one year. Fourth, at least 10% of firm switchers from firm f also move to f' within the year.

The first two conditions ensure that firm f undergoes a large enough change change in its occupational structure at the same time as its spends more on outsourced labor. The third

⁷Our results are virtually identical when varying the number of groups between 10 and 200.

⁸The standard deviation of worker effects is 0.45, while the standard deviation of log wages is 0.54. Thus, worker and firm effects account for 68% and 7% of the variance of log wages respectively.

⁹Notice that from our theory, we cannot conclude that they hire all workers from non-employment. Indeed, depending what tie-breaking rule we adopt to determine whether a worker at wage w who receives another job offer at the same wage w switches employer or not, contractor firms may or may not hire some workers from employment at other contractor firms.

¹⁰We define unemployment in the matched employer-employee data as in Bilal (2020).

Table 3: Wage penalty from outsourcing.

	Post-Outsourcing		Pre-Outsourcing	
	(1)	(2)	(3)	(4)
Outsourcing Event	-0.140*** (0.037)	-0.119** (0.039)	-0.006 (0.006)	-0.005 (0.006)
Employer Switch		-0.026*** (0.007)		-0.001 (0.001)
<i>Fixed Effects</i>				
2-digit occupation	✓	✓	✓	✓
Obs.	874650	874650	1060032	1060032

Note: Standard errors in parenthesis, clustered by 2-digit occupation. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Due to changes in the occupation classification in 2002, we lose the time period 1997-1999 and only have two time periods and thus no time period fixed effects.

and fourth conditions isolate events in which a large enough group of workers transitions to the same destination employer as in Goldschmidt and Schmieder (2017). We use a lower threshold than them to maximize statistical power, as we must work with a 4% panel to track workers across years.

We consider an econometric specification of the form:

$$\Delta^2 \log w_{i,t} = \delta_{o(i)} + \mu \mathcal{O}_{i,t_0(i)} + \nu \mathcal{J}_{i,t_0(i)} + v_{i,t}. \quad (11)$$

t_0 denotes the year of the outsourcing event, $\mathcal{O}_{i,t_0(i)}$ is a dummy variable that equals one if the worker is an outsourcing event in year $t_0(i)$. $\mathcal{J}_{i,t_0(i)}$ denotes an indicator that equals one if worker i changes employer in year $t_0(i)$ to control for the possible common effect of switching employer. $\Delta^2 \log w_{i,t}$ denotes wage growth between year t and year $t_0(i)$. We use the notation Δ^2 to indicate that we also remove worker-specific linear trends in wages that may confound our results. We estimate those trends using only years preceding the outsourcing event $t < t_0$. Finally, $\delta_{o(i)}$ are two-digit occupation fixed effects that capture occupation-specific trends.

Table 3 displays our results. In column (1), we find that an outsourcing event is associated with a 14% wage decline on average over the subsequent four years. Since some of that estimate may be driven by displacement-like effects, we also control for the common effect of switching employers in column (2). Our coefficient of interest declines modestly to 12% and remains statistically significant. Interestingly, it coincides almost exactly with the negative contractor firm wage premium that we estimated in the two-way fixed effect specification (10). Columns

(3-4) show that there is no evidence of pre-outsourcing effects. We conclude that outsourcing is associated to a substantial wage penalty.

To summarize, this section has proposed reduced-form evidence supporting the key predictions of our theory: selection into outsourcing, as well as the productivity and distributional effects. Having validated the core structure of our theory of outsourcing, we now turn to our general equilibrium quantitative exercises.

4 Extended model and estimation

In this section, we start by presenting the extended version of the model from Section 2 that we take to the data. We then discuss our estimation strategy.

4.1 Quantitative setup

Vacancies. Instead of being endowed with a unit measure of vacancies, we now let traditional firms post any number of vacancies v in market s at a convex cost

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

for $\gamma > 0$. When $\gamma \rightarrow +\infty$ we recover the model of Section 2.¹¹

Since not all traditional firms post the same number of vacancies in equilibrium, the number of workers a traditional firm attracts and retains must reflect its vacancy share:

$$n_s(w, v) = \frac{(1 + k_s)e_s}{(1 + k_s(1 - F_s(w)))^2} \cdot \frac{v}{V_s}. \quad (12)$$

where

$$V_s = \mathcal{V}_s + M \sum_{\mathbf{o}} \int v_s(z, \mathbf{o}) \Omega(\mathbf{o}|z) d\Gamma(z) \quad (13)$$

denotes the equilibrium total number of vacancies in market s . $v_s(z, \mathbf{o})$ denotes the equilibrium number of vacancies posted by a firm with productivity z and outsourcing decision $\mathbf{o} = \{o_s\}_s \in \{0, 1\}^S$. $\Omega(\mathbf{o}|z)$ denotes the share of firms with productivity z that choose the outsourcing

¹¹We do not introduce a type-specific shifter in the vacancy cost function because it would simply be a normalization of what “one” vacancy is once we introduce a type-specific matching function efficiency below.

bundle $\mathbf{o}$, which we characterize in equation (16) below.

Equation (12) is the analog of equation (2) when vacancies vary across firms. The first term on the right-hand-side is identical to that in equation (2), having anticipated that the wage offer distribution has no mass points except at the reservation wage. The second term on the right-hand-side replaces $\frac{1}{M}$ in equation (2), and is the firm's vacancy share.

Idiosyncratic outsourcing costs. We introduce firm-specific additive costs of outsourcing, $\varepsilon_{\mathbf{o}}$, that are indexed by the entire vector of outsourcing decisions. We assume that they are independent across firms and outsourcing bundles and follow a Gumbel distribution with shape parameter ι .

Matching function. We assume that a Cobb-Douglas matching function determines the contact rates λ_s^U, λ_s^E through the total number of matches per unit of time

$$\mathcal{M}_s = \mu_s (m_s(u_s + \zeta_s(1 - u_s)))^\xi V_s^{1-\xi}. \quad (14)$$

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

Revenue function, productivity distribution and benefits. We assume that the revenue function takes the form

$$R(z, \{n_s\}_s) = z \left(\int (a_s n_s)^\alpha ds \right)^{\frac{\rho}{\alpha}}.$$

where $\min_s a_s$ is normalized to 1. The elasticity of substitution between worker types is $\frac{1}{1-\alpha}$. This revenue function is consistent with both micro-foundations discussed in examples 1 and 2 in Section 1.3. We assume a lognormal productivity distribution with logged variance κ^2 .¹² Unemployment benefits are $b_s = b \cdot a_s$.

Traditional firms' decision. Traditional firms then solve $\pi(z) = \max_{\{o_s\}_s \in \{0,1\}^S} \pi(z|\{o_s\}_s)$, where:

$$\pi(z|\{o_s\}_s) = \max_{\{n_s\}_s, \{v_s\}_s} R(z, \{n_s\}_s) - \int \left\{ [(1 - o_s)w_s(n_s, v_s) + o_s p_s] n_s + (1 - o_s)c(v_s) \right\} ds + \varepsilon_{\{o_s\}_s}, \quad (15)$$

¹²The logged mean is a simple normalization of the economy, and so is set to 0.

where $n \mapsto w_s(n, v_s)$ is the inverse function of $w \mapsto n_s(w, v_s)$ as defined in equation (12). As a result, the share of firms of productivity z that choose bundle is $\mathbf{o}$ given by

$$\Omega(\mathbf{o}|z) = \frac{e^{\iota\pi(z|\mathbf{o})}}{\sum_{\mathbf{o}'} e^{\iota\pi(z|\mathbf{o}')}}. \quad (16)$$

Wage offer distribution. As in Section 2, more productive firms post higher wages conditional on their outsourcing bundles $\mathbf{o}$. However, because traditional firms with different outsourcing bundles $\mathbf{o}$ may post different wages even at the same productivity z , the equilibrium number of workers that traditional firm z attracts and retains cannot be expressed as a function of $\Gamma(z)$ alone. Instead, it is a function of $\Upsilon_s(z, \mathbf{o}) = F_s(w_s^*(z, \mathbf{o}))$, where the wage offer distribution F_s solves

$$F_s(w) = \sum_{\mathbf{o}} \int \mathbb{1}\{w_s(z, \mathbf{o}) \leq w\} v_s(z, \mathbf{o}) \Omega(\mathbf{o}|z) d\Gamma(z). \quad (17)$$

Number of skill types. We solve the model for two skill types.

Having laid out the structure of the extended framework, we turn to the estimation strategy.

4.2 Estimation strategy

We set a quarterly frequency. We then define skill groups as revealed by the occupations of workers. We rank 2-digit occupations by their average wage. We compute the mean predicted wage of a worker by interacting the mean occupational wages and the time spent by the worker in each occupation. A worker is low-skill if their predicted wage is below median, and high-skill if above. We use non-employment as our primary measure of “unemployment” in the model, since a large fraction of steady-state flows into employment stem from individuals officially out of the labor force.

We then estimate the model in three steps. First, we estimate a first group of parameters that can be straightforwardly mapped to data. Second, we set one parameter to a specified value. In the third step, we estimate the remaining parameters jointly with a MSM estimator.

First, we identify the parameters $\{\delta_s, \zeta_s\}_{s=1}^2$ from labor market flows and the mass of firms M from firm size. The time-aggregated employment-to-non-employment transition rate EN_s is equal to the job losing rate parameter δ_s . In addition, in the extended model of Section 4.1, all matches are viable due to the Inada property of the revenue function. Therefore, all

meetings from non-employment result in a viable match, implying that the time-aggregated non-employment-to-employment transition rate NE_s is equal to the endogenous offer rate λ_s^U from non-employment. We then relate the employment-to-employment transition rate EE_s to underlying arrival rates. We show in Appendix B.1 that

$$\frac{EE_s}{EN_s} = \frac{(1 + k_s) \log(1 + k_s) - k_s}{k_s}.$$

Since $k_s = \lambda_s^E / \delta_s$, this relationship immediately determines the endogenous offer rate from employment, λ_s^E . Thus, we recover $\zeta_s = \lambda_s^E / \lambda_s^U$. Finally, since we will match flows out of non-employment as well, we choose M to match average firm size.

Second, we set the elasticity of the matching function ξ to a pre-specified value. Since our data does not allow us to credibly estimate it, we set to $\xi = 0.5$, a central value found in the literature as reviewed by Petrongolo and Pissarides (2001).

Third, we jointly estimate the remaining parameters $\{\mu_s\}_{s=1}^2, b, a_2, \rho, \alpha, \kappa, \gamma, \{\tau_s\}_{s=1}^2, \iota$ by MSM. While the parameters are of course jointly identified, we provide an heuristic argument that describes how the moments we choose inform parameters. We confirm our argument numerically in Section 4.3 below.

Inspection of equation (14) reveals that the matching function efficiency for skill s , μ_s , has a direct impact on the non-employment-to-employment transition rates NE_s which we target for each skill. By shifting unemployment benefits conditional on wages, the parameter b affects the average replacement rate in the economy which we target. We inform the relative productivity of high-skill workers a_2 using the skill premium. The curvature in the revenue function ρ shifts average profits in the economy conditional on wages, and so we target the aggregate labor share.

The parameter α governs the elasticity of substitution between skills, and thus drives how relative wages respond to relative employment conditional on the wedge that labor market frictions introduce between wages and marginal products of labor as per Proposition 3. Two forces operate. First, the job ladder implies that firms with higher wages have more workers, pushing towards a positive correlation between relative wages and relative employment. Second, the basic substitution effect implies that the relative marginal product of labor is decreasing in relative employment, and more so if α is high. Thus, by targeting the standardized regression coefficient of relative wages onto relative employment of the value added distribution, we inform α .¹³

¹³See for instance Acemoglu and Autor (2011). We found this moment to be most informative about α when restricting attention to the 25th to 50th percentiles of the value added distribution, which we do in the model

Given ρ , γ and κ then jointly determine the dispersion in size and value added. When there is less curvature in the vacancy cost, productive firms are able to hire more and dispersion in firm size increases. Conditional on size, the dispersion in productivity raises dispersion in value added. Thus, we target the standard deviation of log firm size to inform γ , and the ratio between the standard deviation of log value added and the standard deviation of log firm size to inform κ .

The outsourcing cost parameters pin down how much outsourcing there is in the economy. We target aggregate outsourcing expenditures as a fraction of wages in 1996 to pin down the common level of $\{\tau_s\}_{s=1}^2$. We then target the employment share in high-skill contractor firms relative to the employment share in low-skill contractor firms to find τ_2/τ_1 .¹⁴ The dispersion parameter ι governs how much outsourcing occurs at the bottom of the productivity distribution. As per equation (16), the lower ι , the more mixing there is in outsourcing decisions relative to the model in Section 2 without idiosyncratic costs.¹⁵ Thus, we also target the outsourcing expenditure share for firms with below-median value added, a counterpart of the reduced-form coefficient for selection into outsourcing estimated in Section 3.3. With this moment, we discipline the relationship between productivity and outsourcing which drives the productivity effect.

We use a loss function with squared proportional deviations for each moment, weight moments equally and use a gradient descent algorithm to find the minimum.

4.3 Estimation results and identification

Table 4 summarizes our estimation results. Our parameter estimates fall within conventional ranges found in the literature. The revenue function curvature parameter $\hat{\rho} = 0.90$ translates into an elasticity of substitution between varieties of about 10, a value towards the upper end of values found in the literature. Our substitutability parameter $\hat{\alpha} = 0.180$ corresponds to an elasticity of substitution between workers of 1.21, close to typical estimates. The curvature in the vacancy cost $\hat{\gamma} = 7.01$ lies towards the upper end of usual values.

Our estimate of outsourcing costs $\hat{\tau}_1 = 0.49$ implies that outsourcing low-skill labor implies a

and the data.

¹⁴While the employment share in contractor firms identified by industry codes for any given skill type may be subject to some measurement error, using the relative employment shares and the aggregate expenditure share may provide more accurate targets.

¹⁵Formally, with a large enough number of skills, it is straightforward to see that $\Omega(\mathbf{o}|z)$ is log-supermodular in (ι, z) .

Table 4: Parameter estimates and empirical targets.

Parameter	Interpretation	Target	Empirical Moment	Simulated Moment	Parameter Estimate
A. Parameters from direct inversion.					
δ_1	Job loss rate low-skill	EN rate low-skill	0.04		0.04
δ_2	Job loss rate high-skill	EN rate low-skill	0.03		0.03
ζ_1	Rel. search. emp. low-skill	EE rate low-skill	0.04		0.75
ζ_2	Rel. search. emp. high-skill	EE rate low-skill	0.03		0.68
M	Measure of firms	Average firm size	8.10		0.10
B. Parameters from MSM estimator.					
μ_1	Matching eff. low-skill	NE rate low-skill	0.17	0.17	0.42
μ_2	Matching eff. high-skill	NE rate high-skill	0.17	0.18	0.43
a_2	Rel. prod. high-skill	Skill premium	1.74	1.73	1.85
b	Unemployment benefits	Replacement rate	0.70	0.70	55.0
ρ	Curvature in revenue	Labor share	0.70	0.71	0.90
α	Subst. between skills	Skill prem. by rel. emp.	0.70	0.69	0.18
γ	Curvature vac. cost	St. dev. log firm size	0.98	0.98	7.01
κ	Standard dev. prod.	St.d. V.A. / St.d. size	1.16	1.17	0.41
τ_1	Out. cost low-skill	Out. share V.A.>P50	0.06	0.06	0.49
τ_2	Out. cost high-skill	Rel. emp. low-skill out.	2.50	2.49	0.28
ι	Dispersion out. costs	Out. share V.A.<P50	0.04	0.04	0.02

markup around 2 over the wages paid to the workers providing the labor services. The estimate $\hat{\tau}_2 = 0.28$ implies that outsourcing high-skill labor requires a markup around 4, reflecting the relative scarcity of outsourced high-skill workers. To the best of our knowledge, we are the first to propose such estimates, and thus have no benchmark to compare them to.

How well are parameters identified? To answer this question, Figure 3 plots both the simulated moment and the loss function as we vary the parameter close to its estimated value. If the simulated moment line is steep, then the parameter is locally well identified in a univariate sense. If the loss function curve is highly peaked, then the parameter is locally well identified in a multivariate sense. Lest we simulate the model on the full twelve-dimensional hypercube however, we cannot guarantee global identification.

Figure 3 reveals that most of the parameters are well identified, which can be seen from the moment deviation being steep as a function of parameter deviations, and the loss function

Figure 3: Simulated moments and loss function across parameter values.

being peaked around 0. The curvature of the vacancy cost γ is not as sharply identified as other parameters, neither from a univariate nor multivariate perspective. The dispersion parameter for idiosyncratic outsourcing costs ι is well identified in a univariate sense, but not as sharply in a multivariate sense. Yet, overall, Figure 3 confirms our identification argument and gives us confidence in the credibility of our estimation strategy.

Having discussed the mapping between parameters and targeted moments, we now turn to our quantitative results.

5 The effects of outsourcing on inequality and output

In this section, we first presents how the estimated model fares relative to several non-targeted moments that relate to the distributional effect. Second, we present our general equilibrium results describing the aggregate effects of outsourcing.

5.1 Over-identification and the distributional effect

Our estimation strategy left out several key moments that relate to the distributional effect of outsourcing. We now check the validity of the model by assessing whether it can match those.

Wage dispersion. The first moment that is central to the distributional effect of outsourcing is the steady-state dispersion in firm wage premia. Indeed, as per Proposition 4, it governs the wage loss of a worker who falls from the top to the bottom of the job ladder. Importantly, within-skill dispersion in wages is primarily driven by how the model maps labor market flows $k_s = \lambda_s^E / \delta_s$ and the distribution of marginal products of labor into wages. Thus, our estimation strategy has imposed few direct restrictions on wage dispersion.

To estimate firm wage premia consistently with our theory, we run equation (10) similarly to Section 3.5 but separately by skill group. We estimate the within-skill standard deviation of firm wage premia to be 0.14 in the data. In our estimated model, the within-skill standard deviation of log wage premia is 0.15. We conclude that the estimated model accounts for observed wage dispersion.

Outsourcing wage penalty. We now turn to the outsourcing wage penalty. Our estimation strategy only targets expenditures on outsourcing by traditional firms, and does not restrict the wage gap between traditional and contractor firms.

Using an event study, we highlighted in Section 3.5 that outsourced workers lose on average 12 log points relative to pre-outsourcing wages. We replicate the event study in the estimated model. Doing so necessarily involves additional assumptions. We consider traditional firms that start in an initial period with a set of idiosyncratic shocks for outsourcing bundles $\{\varepsilon_{\mathbf{o}}\}_{\mathbf{o}}$. We then assume that they draw a new set of idiosyncratic shocks $\{\varepsilon'_{\mathbf{o}}\}_{\mathbf{o}}$ in a second period, and choose their preferred production structure anew. Importantly, we assume that the new shocks $\{\varepsilon'_{\mathbf{o}}\}_{\mathbf{o}}$ are independent from their initial draw $\{\varepsilon_{\mathbf{o}}\}_{\mathbf{o}}$.¹⁶ Depending on the new shocks, traditional firms may lay off workers of any skill to rent labor services from contractor firms. Finally, we assume that the mass of firms who re-draw these shocks is small enough that workers who lose their job due to an outsourcing event transition into contractor firms.

Armed with those assumptions, we compute the outsourcing wage penalty in the estimated model. We find that outsourced workers' wages drop by 17 log points on average. While

¹⁶The joint probability that a firm with productivity z chooses the sequence of bundles $\mathbf{o}, \mathbf{o}'$ is then given by $\Omega(\mathbf{o}|z)\Omega(\mathbf{o}'|z)$.

this value overstates the point estimate from Section 3.5 by 5 log points, it lies within the 95% confidence interval and thus we cannot reject that they are equal. We conclude that the estimated model accounts for the outsourcing wage penalty.

Outsourced employment share. Whether outsourcing is of macroeconomic relevance depends not only on outsourcing expenditures, but also on how many workers work at contractor firms. In our estimation, we targeted aggregate outsourcing expenditures in 1996 and the relative employment shares at contractor firms, but did not restrict the level of employment at contractor firms.

We showed in Section 3.2 that contractor firms employ 5% of low skill workers and 2% of high-skill workers at the beginning of our sample. The estimated model predicts that 6% of low skill and 2% of high-skill workers are employed by contractor firms. We conclude that the estimated model accounts for outsourced employment.

Having shown that the estimated model replicates key non-targeted moments, we turn to our main counterfactual exercise.

5.2 The aggregate effects of outsourcing

We investigate the effect of the rise in outsourcing on inequality and aggregate output. To do so, we lower outsourcing costs $(1/\tau_1, 1/\tau_2)$ by the same proportional amount, in order to replicate the increase in the aggregate expenditure share on outsourcing from Figure 1(a) between 1996 and 2007. We compare steady-states of the estimated model, and interpret our results as the effect of outsourcing on the French labor market in that decade. Figure 4 displays the results. All outcomes are shown as a function of aggregate outsourcing expenditures relative to the aggregate wage bill on the x-axis.

Panel (a) shows the fraction of workers employed at contractor firms. As the cost of outsourcing falls, the fraction of low-skill workers at contractor firms rises from 6% to almost 10%, while that of high-skill workers increases from 2% to 3%. Given the outsourcing wage penalty, this reallocation of workers towards contractor firms has a negative partial equilibrium impact on workers' average earnings.

Panels (b) and (c) reveal that general equilibrium effects further deteriorate workers' prospects. Wages of both low-skill and high-skill workers fall across the distribution, due to two effects. First, the fall in the price of outsourcing implies a direct compression effect on top wages high-

Figure 4: The effects of outsourcing on inequality and output.

lighted in Proposition 4.4. Second, the reallocation of workers away from high-paying firms and towards contractor firms at the bottom of the job ladder weakens labor market competition for workers from both ends of the job ladder. Middle-productivity firms that hire in-house are now shielded from losing their workers to a high-productivity firm with more generous wage offers. In addition, middle-productivity firms can now poach more workers from low-paying contractor firms. Together, these observations imply that middle-productivity firms lower pay and still recruit as many workers as they need. This effect is particularly potent for in-house low-skill workers around the median of the wage distribution, who lose almost 5% in wages. As a result of the general equilibrium decline in wages throughout the distribution, the reservation wage also falls and workers at the bottom decile also lose, albeit more modestly.

Outsourcing has not only within-skill distributive effects, it also affects between-skill inequality. Because low-skill workers are relatively more affected by outsourcing than high-skill workers as seen in panels (b) and (c), the skill premium rises. Panel (d) shows however that

this increase is modest, less than one log point.

Together, panels (a) through (d) reveal that outsourcing deteriorates the prospects of workers in the labor market. Panel (e) confirms that the distributional effect of outsourcing dominates the productivity effect for workers. Aggregate labor income falls by almost 2% as outsourcing rises between 1996 and the counterfactual 2007 economy. This reduction masks two offsetting effects. Because it improves the aggregate allocation of efficiency units of labor, outsourcing leads to more hiring in the aggregate, and the non-employment rate drops by 1.5 percentage points. However, the decline in wages conditional on work highlighted in panels (b) and (c)—the intensive margin—dominates the non-employment gains—the extensive margin.

Outsourcing still leads to a productivity effect, and aggregate output rises. Aggregate output gains are, however, modest: less than 1% between 1996 and the counterfactual 2007 economy. This small increase in output stems from two competing effects. First, the reduction in outsourcing costs leads more firms to outsource and scale up, boosting aggregate output. However, this reallocation of workers requires contractor firms to spend more resources on recruiting. Workers at contractor firms are frequently poached by better-paying in-house traditional firms because contractor firms are at the bottom of the job ladder. As a result, contractor firms end up posting a large amount of vacancies: their vacancy share rises by 16 percentage points between 1996 and the counterfactual 2007 economy. This increase in the market share of contractor firms makes hiring more difficult for traditional firms who remain in-house. Although they pay lower wages, they also cut back on hiring and shrink, adversely affecting aggregate output.

Combining these aggregate impacts, panel (f) shows that outsourcing strongly depresses the labor share by more than 3 percentage points between 1996 and the counterfactual 2007 economy. We conclude that outsourcing has modest positive productivity effects, but these gains benefit firm owners at the expense of workers.

Conclusion

This paper started with a theory of domestic outsourcing. We have argued that it is useful to conceptualize firms' outsourcing decision in the context of frictional labor markets giving rise to firm wage premia. More productive firms are then more likely to outsource. Outsourcing raises output at the firm level. And labor service providers endogenously locate at the bottom of the job ladder, implying that outsourced workers receive lower wages. Together, these observa-

tions characterize the tension between productivity enhancements and redistribution away from workers that is tied to outsourcing. Using firm-level instruments for outsourcing and revenue productivity, we have proposed new reduced-form evidence that confirms the productivity and redistributive effects of outsourcing. Finally, equipped with a structurally estimated model, we have shown that outsourcing benefits firm owners and deteriorate workers' prospects in the aggregate.

There are at least two natural directions along which to expand this research agenda. First, the productivity and distributional effects of outsourcing could be more fully contrasted with those from trade in intermediate goods and services than the stark perspective we have taken to highlight the unique features of outsourcing. Second, due to its tractability under parsimonious assumptions, our framework is naturally equipped to study questions with an efficiency-equity trade-off that involve workers' wages and scale-biased aggregate transformations, such as trade liberalizations 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.
- Acemoglu, Daron and Autor, David (2011). “Skills, Tasks and Technologies: Implications for Employment and Earnings”. In: *Handbook of Labor Economics, Volume 4. Amsterdam: Elsevier-North*, pp. 1043–1171.
- Acemoglu, Daron, Gancia, Gino, and Zilibotti, Fabrizio (2015). “Offshoring and Directed Technical Change”. *American Economic Journal: Macroeconomics* 7.3, pp. 84–122.
- Antràs, Pol, Fort, Teresa C., and Tintelnot, Felix (2017). “The Margins of Global Sourcing: Theory and Evidence from U.S. Firms”. *American Economic Review* 107.9, pp. 2514–64.
- Bergeaud, Antonin, Malgouyres, Clement, Mazet-Sonilhac, Clement, and Signorelli, Sara (2020). *Technological Change and Domestic Outsourcing*. Working Paper.
- Bertrand, Marianne, Hsieh, Chang-Tai, and Tsivanidis, Nick (2020). *Contract Labor and Firm Growth in India*. Working Paper.
- Bilal, Adrien (2020). *The Geography of Unemployment*. Working Paper.
- Bonhomme, Stéphane, Lamadon, Thibaut, and Manresa, Elena (2019). “A Distributional Framework for Matched Employer Employee Data”. *Econometrica* 87.3, pp. 699–739.
- Burdakov, Oleg, Dai, Yu-Hong, and Huang, Na (2019). *Stabilized Barzilai-Borwein method*. Tech. rep.
- Burdett, Kenneth and Mortensen, Dale T. (1998). “Wage Differentials, Employer Size, and Unemployment”. *International Economic Review* 39.2, pp. 257–273.
- Cahuc, Pierre, Postel-Vinay, Fabien, and Robin, Jean-Marc (2006). “Wage Bargaining with On-the-Job Search: Theory and Evidence”. *Econometrica* 74.2, pp. 323–364.
- 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 (May 2013). “Workplace Heterogeneity and the Rise of West German Wage Inequality”. *The Quarterly Journal of Economics* 128.3, pp. 967–1015.

- Drenik, Andres, Jäger, Simon, Plotkin, Miguel Pascuel, and Schoefer, Benjamin (2020). *Paying Outsourced Labor: Direct Evidence from Linked Temp Agency-Worker-Client Data*. Working Paper 26891. National Bureau of Economic Research.
- Engbom, Niklas and Moser, Christian (2018). *Earnings Inequality and the Minimum Wage: Evidence from Brazil*. CESifo Working Paper Series 6393. CESifo Group Munich.
- 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.
- Goldschmidt, Deborah and Schmieder, Johannes F. (Apr. 2017). “The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure”. *The Quarterly Journal of Economics* 132.3, pp. 1165–1217.
- Gouin-Bonenfant, Emilien (2018). *Productivity Dispersion, Between-firm Competition and the Labor Share*. 2018 Meeting Papers 1171. Society for Economic Dynamics.
- Grossman, Gene M. and Rossi-Hansberg, Esteban (2008). “Trading Tasks: A Simple Theory of Offshoring”. *American Economic Review* 98.5, pp. 1978–97.
- Helpman, Elhanan, Itskhoki, Oleg, and Redding, Stephen (2010). “Inequality and Unemployment in a Global Economy”. *Econometrica* 78.4, pp. 1239–1283.
- Hummels, David, Jorgensen, Rasmus, Munch, Jakob, and Xiang, Chong (2014). “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data”. *American Economic Review* 104.6, pp. 1597–1629.
- Katz, Lawrence F. and Krueger, Alan B. (2017). “The Role of Unemployment in the Rise in Alternative Work Arrangements”. *American Economic Review* 107.5, pp. 388–92.
- LeMoigne, Mathilde (2020). *Exploring the ‘Fissured workplace’: Internal job ladders’ fragmentation and its effect on plants and workers*. Working Paper.
- Petrongolo, Barbara and Pissarides, Christopher A. (2001). “Looking into the Black Box: A Survey of the Matching Function”. *Journal of Economic Literature* 39.2, pp. 390–431.
- Topkis, Donald M. (1998). *Supermodularity and Complementarity*. Princeton University Press. ISBN: 9780691032443.

A Proofs

A.1 Reservation wage

Omit s indices. Suppose without loss of generality that F admits a density f . Then,

$$\left[r + \delta + \lambda^E(1 - F(w)) \right] V(w) = w + \delta U + \lambda^E \int_w^\infty V(x) f(x) dx.$$

Differentiate w.r.t. w to obtain

$$\left[r + \delta + \lambda^E(1 - F(w)) \right] V'(w) = 1.$$

Integrate back to

$$V(w) = U + \int_{\underline{w}}^w \frac{dx}{r + \delta + \lambda^E(1 - F(x))}.$$

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 $V(\underline{w}) = U$,

$$(r + \lambda^U)U = b + \lambda^U \int_{\underline{w}}^\infty V(x) f(x) dx, \quad (r + \lambda^E)U = \underline{w} + \lambda^E \int_{\underline{w}}^\infty V(x) f(x) dx,$$

so that

$$rU = \frac{\lambda^U \underline{w} - \lambda^E b}{\lambda^U - \lambda^E}.$$

Therefore,

$$\lambda^U \underline{w} = \lambda^E b + (\lambda^U - \lambda^E) \left[b + \lambda^U \int_{\underline{w}}^\infty \frac{(1 - F(x))dx}{r + \delta + \lambda^E(1 - F(x))} \right] \quad (18)$$

A.2 Proof of equation (1)

The flow of workers out of any wage interval $[\underline{w}_s, w)$ must be equal to the flow of workers into that wage interval:

$$\lambda_s^U F_s(w) u_s = \left(\delta_s + \lambda_s^E (1 - F_s(w)) \right) (m_s - u_s) G_s(w),$$

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 the expression in equation (1).

A.3 Dynamic firm problem

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

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

where $\phi = \frac{u}{u + \frac{\lambda^E}{\lambda^U}(1-u)}$ is the probability of meeting an unemployed worker. Re-arrange to obtain

$$\phi = \frac{1}{1 + k}. \quad (19)$$

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

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

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

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

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

Using (19),

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

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

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

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

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

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

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

Substituting into the first-order condition, we obtain

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

which coincides with the static first-order condition.

A.4 Proof of Proposition 1

Consider the profit function

$$\Pi[z, \{w_s\}_s] = R[z, \{n_s(w_s)\}_s] - \int w_s n_s(w_s) ds$$

Since n_s is increasing in w , Π is 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 2.8.1. p. 76 in Topkis (1998). Thus, the set of maximizers $\{w_s(z)\}_s$ are strictly increasing in z for each s . The argument to rule out mass points and connected support then follows directly Burdett and Mortensen (1998).

A.5 Proof of Proposition 2

Given the ordering of the ordering of wages from Proposition 1, $F(w_s(z)) = \Gamma(z)$, and $n_s(z) = \frac{(1+k_s)e_s}{M[1+k_s(1-\Gamma(z))]}$.

A.6 Proof of Proposition 3

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

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

where $\mathbf{n}_s$ denotes the vector $\mathbf{n}$ without its entry s . Multiplying both sides by $w'_s(z)$, integrating over z and changing variables to $n_s(w_s(z)) \equiv n_s(z)$ delivers the formula in Proposition 3.

A.7 Micro-foundations for the cost of outsourcing

1. Iceberg trade cost. To sell one unit of labor services to a traditional firm, contractor firms must hire $1/\tau_s$ units of labor.

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

problem of the contractor firm is

$$\pi^C(w) = \max_k p_s k^{1-\beta} n_s(w)^\beta - r_s k - w n_s(w). \quad (20)$$

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 (20), we obtain $\pi^C(w) = p_s \left(\frac{K_s}{Q_s^{\text{Out}}} \right)^{1-\beta} n_s(w) - w n_s(w)$. Assume further that $K_s = \tau_s^{\frac{1}{1-\beta}}$, and take $\beta \rightarrow 1$. Then, (20) becomes

$$\pi^C(w) = (\tau_s p_s - w) n_s(w).$$

A.8 Proof of Proposition 4

1. Contractor firm wage. The free-entry condition (4) immediately implies that contractor firms pay the reservation wage. If they posted $w > \underline{w}_s$, they could deviate to $w - \varepsilon$ for a small ε and make positive profits, a contradiction.

2. Price of outsourcing. The price of outsourcing also follows immediately from the free-entry condition (4) when there is some outsourcing in equilibrium.

3. Outsourcing threshold. We immediately see that the profit function in (5) is super-modular in $(z, \{n_s\}_s, \{o_s\}_s)$. We again use Theorem 2.8.1. p. 76 in Topkis (1998) to obtain that wages and size are rising in productivity. In addition, we also obtain that the outsourcing decision is rising in productivity. Since it is binary, there must be a threshold productivity $\hat{z}_s$ such that firms outsource if and only if $z \geq \hat{z}_s$.

4. Wage cap. Cost-minimization immediately implies that $p_s \geq w_s(\hat{z}_s)$.

5. Size and revenue jump and strict wage cap. For notational simplicity and without loss of generality, focus on the case with a single worker type in the remainder of this proof and drop s indices. Denote outsourced employment by q .

Size jump. We start by showing that $q(\hat{z}) > n(\hat{z})$.

Proof. Because of the theorem of the maximum, profits must be continuous at the outsourcing cutoff $\hat{z}$. Then, the indifference condition at $\hat{z}$ becomes

$$R(\hat{z}, n(\hat{z})) - w(n(\hat{z}))n(\hat{z}) = R(\hat{z}, q(\hat{z})) - pq(\hat{z})$$

where $w(n)$ is an increasing function. In addition, the first-order condition for in-house employment is $R_n(\hat{z}, n(\hat{z})) = w(\hat{z}) + n(\hat{z})w'(n(\hat{z}))$ for in-house production. For outsourcing it is $R_n(\hat{z}, q(\hat{z})) = p$. Substituting both into the indifference condition (21):

$$R(\hat{z}, n(\hat{z})) - R_n(\hat{z}, n(\hat{z}))n(\hat{z}) + n(\hat{z})^2w'(n(\hat{z})) = R(\hat{z}, q(\hat{z})) - pR_n(\hat{z}, q(\hat{z}))$$

Since we assumed that R be strictly concave in n , the function $R - nR_n$ is strictly increasing in n . To finish the proof, we thus need to show that $w'(n(\hat{z})) > 0$. To that end, notice that $n'(w(\hat{z})) = 2(1+k)ekF'(w(\hat{z}))$. From the ordering of wages, $F'(w(\hat{z}))w'(\hat{z}) = \Gamma'(\hat{z}) < +\infty$. Thus, $F'(w(\hat{z})) < +\infty$, implying that $n'(\bar{w}) < +\infty$. As a result, $w'(n(\hat{z})) = \frac{1}{n'(\bar{w}(\hat{z}))} > 0$. $\square$

Revenue jump. Since labor inputs jump up at the outsourcing cutoff, so does revenue.

Strict wage cap. Since the marginal in-house producer could have chosen to outsource the same-sized workforce, it must be that $R(\hat{z}, n(\hat{z})) - w(n(\hat{z}))n(\hat{z}) \geq R(\hat{z}, n(\hat{z})) - pn(\hat{z})$. Given that $n(\hat{z}) < q(\hat{z})$, we obtain $p > w(n(\hat{z}))$.

B Simulation and estimation

B.1 Expression for the employment-employment transition rate

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

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

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

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

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

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

B.2 Numerical solution

The number of workers of skill s a traditional firm with productivity z hires depend on the outsourcing bundle, $\mathbf{o}$. If the skill is outsourced, $o_s = 1$, demand solves

$$\rho z a_s n_s(z, \mathbf{o})^{\alpha-1} \left(\sum_t a_t n_t(z, \mathbf{o})^\alpha \right)^{\frac{\rho}{\alpha}-1} = \frac{w_s}{\tau_s}. \quad (21)$$

If the skill is hired in-house, $\mathbf{o}_s = 0$, the number of workers is jointly determined by the position of the firm in the ladder, (12), and the number of vacancy posted, given by

$$\left[\rho z a_s n_s(z, \mathbf{o})^{\alpha-1} \left(\sum_t a_t n_t(z, \mathbf{o})^\alpha \right)^{\frac{\rho}{\alpha}-1} - w_s(z, \mathbf{o}) \right] \tilde{n}_s(z, \mathbf{o}) = V_s^{1+\gamma} \tilde{v}_s(z, \mathbf{o})^\gamma, \quad (22)$$

with $n_s = \tilde{n}_s \cdot \tilde{v}_s$, and $\tilde{v}_s = \frac{v_s}{V_s}$ is the vacancy share. When a traditional firm hires in-house, it also decides on the wage attached to the vacancies, given by the differential equation

$$\partial_z w_s(z, \mathbf{o}) = \frac{2k_s \partial_z \Upsilon_s(z, \mathbf{o})}{1 + k_s \bar{\Upsilon}_s(z, \mathbf{o})} \left[\rho z a_s n_s(z, \mathbf{o})^{\alpha-1} \left(\sum_t a_t n_t(z, \mathbf{o})^\alpha \right)^{\frac{\rho}{\alpha}-1} - w_s(z, \mathbf{o}) \right]. \quad (23)$$

where the notation $\partial_z f(z)$ is a shorthand for $\frac{\partial f}{\partial z}(z)$ for any function f , and $\bar{\Upsilon}_s(z, \mathbf{o}) = 1 - \Upsilon_s(z, \mathbf{o})$.

These three equations, combined with firm size net of vacancy posting, (12), summarize the problem of the firm. The policy functions depend on three aggregate states: the distribution of firms within bundle-skill pair, $\Upsilon_s(\cdot, \mathbf{o})$, the vacancies posted by traditional firms, and the vacancies posted by contractor firms

B.3 Algorithm

To solve numerically for the policy functions and the aggregate states, we discretize the productivity space into 1,000 linearly spaced grid points, denoted $\{z_i\}_{i=1}^{1000}$.

Policy functions When both skills are outsourced, $v_s(z, \mathbf{o}) = 0$, and we find firm size by solving (21) jointly for both skills. When at least one skill is in-house, we solve for the policy functions by iterating upward on the productivity grid starting from $z = \underline{z}$. For $z = \underline{z}$, we know that $w_s(z, \mathbf{o}) = \underline{w}_s$, and we therefore find $\{v_s(\underline{z}, \mathbf{o})\}_s$ by solving (22). For $z_{i+1} > \underline{z}$, we approximate (23) by backward finite difference,

$$w_{i+1}^{\mathbf{o},s} = w_i^{\mathbf{o},s} + (\mathcal{R}_{i+1}^{\mathbf{o},s} - w_i^{\mathbf{o},s}) \frac{v_{i+1}^{\mathbf{o},s} T_{i+1}^{\mathbf{o},s}}{1 + k^s \tilde{\Upsilon}_{i+1}^{\mathbf{o},s}}, \quad (24)$$

where $\mathcal{R}_i^{\mathbf{o},s} = \rho z a_s n_s(z, \mathbf{o})^{\alpha-1} (\sum_t a_t n_t(z, \mathbf{o})^\alpha)^{\frac{\rho}{\alpha}-1}$ is the marginal product of labor of skill s at firm z_i in bundle $\mathbf{o}$, and $T_{i+1}^{\mathbf{o},s} = 2k^s M \Omega_{i+1}^{\mathbf{o}} d\Gamma_{i+1}$. Plugging the approximated differential equation into the vacancy first order condition obtains

$$\tilde{\mathcal{R}}_{i+1}^{\mathbf{o},s} (v_{i+1}^{\mathbf{o},s})^{\alpha-1} = w_i^{\mathbf{o},s} + (\mathcal{R}_{i+1}^{\mathbf{o},s} - w_i^{\mathbf{o},s}) \frac{v_{i+1}^{\mathbf{o},s} d\tilde{\Upsilon}_{i+1}^{\mathbf{o},s}}{1 + k^s \tilde{\Upsilon}_{i+1}^{\mathbf{o},s}} + \frac{(V^s)^{1+\gamma} (v_{i+1}^{\mathbf{o},s})^\gamma}{n_{i+1}^{\mathbf{o},s}},$$

where $\tilde{\mathcal{R}}_{i+1}^{\mathbf{o},s} = \rho z a_s \tilde{n}_s(z, \mathbf{o})^{\alpha-1} (\sum_t a_t n_t(z, \mathbf{o})^\alpha)^{\frac{\rho}{\alpha}-1}$. In principle, we must solve jointly this multidimensional system for vacancies because the marginal product of labor is a function of the vector $\{v_i^{\mathbf{o},s}\}_s$. To avoid the instability of a multidimensional root solver and maintain efficient computation speed, we instead make use of the approximation $\mathcal{R}_{i+1}^{\mathbf{o},s} \approx \mathcal{R}_i^{\mathbf{o},s}$ and $\tilde{\mathcal{R}}_{i+1}^{\mathbf{o},s} \approx \tilde{\mathcal{R}}_i^{\mathbf{o},s}$ which is valid when the grid is fine enough. With these substitutions, we decouple the system and can solve for $\{v_{i+1}^{\mathbf{o},s}\}_s$ independently for each skill while.

This formulation also has the advantage to embed the general equilibrium response of vacancies to changes in wages directly into the equations we solve for. Thus, we sidestep the need to numerically search for a fixed point. Once we compute vacancies, we then obtain firm size and wages via (24). As a result, we only need one iteration across the productivity distribution to solve for policy functions.

Firm distributions Given the firm policy functions, we compute firm's profits and the productivity-specific share of firms per bundle, $\Omega(z, \mathbf{o})$, according to (16). We then create a skill-specific wage grid with endogenous support $[\min_{z, \mathbf{o}} w_s(z, \mathbf{o}), \max_{z, \mathbf{o}} w_s(z, \mathbf{o})]$, and solve for

the skill-specific wage offer distribution, $F_s(\cdot)$, on that grid according to (17). Given F_s , we update the reservation wage according to (18). Finally, using $\Upsilon_s(z, \mathbf{o}) = F_s(w_s(z, \mathbf{o}))$, we recover the within bundle-skill firm distribution by inverting F_s from the wage grid to the productivity grid.

Aggregate vacancies Aggregate vacancies are given by (13). Vacancies posted by service providers are obtained from the market clearing condition for labor services,

$$M \sum_{\mathbf{o}} o_s \int N_s(z, \mathbf{o}) \Omega(z, \mathbf{o}) d\Gamma(z) = \frac{e_s}{1 + k_s} \frac{\mathcal{V}_s}{\mathcal{V}_s + V_s^{in}},$$

the equivalent to (7) in the quantitative model, where V_s^{in} denote aggregate vacancies posted by in-house traditional firms. Both integrals are approximated by the trapezoidal rule.

The code, implemented in Julia, is made of three loops. The inner loop solves jointly the firms' problem and the distribution of firms within bundle-skill pair. The middle loop iterates on vacancies posted by final producers, updating simultaneously the contact rates according to the matching function (14). The outer loop solves for the vacancies posted by service providers.

B.4 Estimation

To estimate the model, we define the loss function

$$\mathcal{L}(\boldsymbol{\theta}) = \sqrt{\frac{1}{M_m} \sum_{m=1}^{M_m} \left[h_m(\boldsymbol{\theta}) - \hat{h}_m \right]^2},$$

where $\boldsymbol{\theta}$ is the vector of parameter to be estimated, $\{\hat{h}_m\}_m$ is the set of empirical moments we are targeting, and $h : \mathbb{R}^{M_m} \mapsto \mathbb{R}^{M_m}$ maps parameters into simulated moments from our model. The simulated moments are computed as exact analogs of the empirical 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 $\{\boldsymbol{\theta}^n\}_n$ by iterating on $\boldsymbol{\theta}^{n+1} = \boldsymbol{\theta} - \gamma_n \nabla \mathcal{L}(\boldsymbol{\theta})$, where the endogenous step size follows the Barzilai–Borwein method. Namely, for $n > 1$,

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

We impose a maximal step size as in Burdakov et al. (2019) to stabilize the descent. The gradient of the loss function is approximated with central finite difference to maximize accuracy. Given that we use $M_m = 11$ parameters, the loss function $\mathcal{L}(\theta)$ is high-dimensional and we cannot check for the existence of local minima. To avoid those, we first search manually to start the algorithm from a θ^0 with a relatively low loss function, in practice $\mathcal{L}(\theta^0) \in [0.3, 0.5]$. The gradient descent then attains the minimum at $\mathcal{L}(\theta^*) = 0.014$. When the gradient is parallelized on 6 CPUs and the descent is run on a standard laptop, the descent takes about one hour.

C Data description

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

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

DADS panel. We use the 4% sample of the DADS panel, between 1996 and 2007. Once a worker enters the dataset in any year after 1976, all her subsequent employment spells are recorded. Individuals’ employment history is recorded in the dataset if (a) they have at least one employment spell, and (b) they are born in October in even years. The dataset provides start and end days of each employment spell, the job’s wage, four-digit occupation and industry, as well as establishment and firm tax identifiers that can be linked to other datasets. We follow Bilal (2020) 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.