

How Does Firm Performance Affect Wages? Evidence from Idiosyncratic Export Shocks*

Andrew Garin Filipe Silverio
Harvard University Banco de Portugal

November 19, 2017

JOB MARKET PAPER

[Please click here for the most recent version](#)

Abstract

In the canonical competitive labor market model, firms are wage-takers and idiosyncratic shocks to individual firms do not affect wages. However, when labor markets are frictional, wages may directly depend on firm-specific factors. We test how sensitive wages are to firm-level labor demand by estimating the incidence of idiosyncratic export demand shocks on the wages of incumbent workers in Portugal during the Great Recession (2008-2010). Using detailed export records, we construct measures of firm exposure to unanticipated shocks to the demands of different countries for specific products. The shocks predict changes in output and payroll at affected firms, but not at other similar firms. We combine the export demand measures with firm balance sheet data and matched longitudinal administrative employer-employee records to estimate the impact of idiosyncratic firm-level demand shocks on employee outcomes. We find that idiosyncratic shocks that decreased sales or value added by 10 percent caused wages of incumbent workers who were employed by affected firms in 2007 to decrease by 1.5 percent relative to trend. Furthermore, we find that these pass-through effects are stronger in industries with higher durability of employment relationships and lower employee turnover rates. These results support a model in which barriers to replacing incumbent workers give rise to internal labor markets within the firm, exposing workers to their employers' idiosyncratic conditions.

*Contact: agarin@fas.harvard.edu. Garin is particularly thankful to Edward Glaeser, Claudia Goldin, Nathaniel Hendren, and Lawrence Katz for extensive advice and support. For thoughtful discussions and comments, we thank Joshua Abel, Lisa Abraham, Pol Antras, Antonio Antunes, David Autor, Alex Bartik, Alex Bell, Allison Bernstein, Laura Blattner, Valentin Bolotnyy, Kirill Borusyak, Sydnee Caldwell, Gabriel Chodorow-Reich, John Coglianesi, David Cutler, Oren Danieli, Yuxiao Huang, Elhanan Helpman, Xavier Jaravel, Dimitri Koustas, Pat Kline, Danial Lakshari, Jessica Liu, Luca Maini, Elizabeth Mishkin, Marc Melitz, Magne Mogstad, Luca Oromolla, Amanda Pallais, Frank Pisch, Pedro Portugal, Heather Sarsons, Alessandro Sforza, Daniel Shoag, Linh Tô, Gabriel Unger, Heidi Williams, Chenzi Xu, and Oren Ziv, as well as seminar participants at Harvard and the Banco de Portugal. All data work was conducted on site at the Banco de Portugal in Lisbon. We are grateful to our colleagues at the Banco de Portugal for making this work possible—especially Amelia Andrade, Antonio Antunes, Isabel Horta Correia, Luisa Farinha, Manuela Lourenco, Pedro Moreira, Francisca Rebelo, and Fatima Teodoro. The views expressed here are ours and do not necessarily reflect those of the Banco de Portugal. We gratefully acknowledge financial support from the Lab for Economic Analysis and Policy at Harvard University and the Multidisciplinary Program in Inequality and Social Policy at the Harvard Kennedy School.

1 Introduction

In the canonical competitive labor market model, wages reflect the supply and demand for skill, and price-taking firms can adjust employment but not wages. However, in practice, labor markets feature barriers that make it costly for firms to replace current employees with new hires and for workers to find comparable new jobs. When firms’ internal labor markets are segmented from the external labor market, the wages of incumbent workers may directly reflect labor demand within the firm apart from the demand level in the external labor market.¹ Recent work documenting substantial wage differentials across firms conditional on fixed worker attributes (Abowd et al., 1999; Card et al., 2016b; Song et al., 2015; Barth et al., 2016) has raised the possibility that heterogeneity in conditions across firms directly contributes to observed wage dispersion—even among workers with the same skills and abilities. This has led some to postulate that noncompetitive rents may play an important role in rising wage inequality, which might be a direct consequence of rising dispersion of firm productivity reflected in unequal worker rents.

Yet, to date, little is known about how much firm-level labor demand can affect wages apart from demand levels in the broader labor market. We present a theoretical framework that illustrates that, in general, one cannot identify the incidence of firm-specific demand changes on wages based on comparisons of observed changes in firm performance and observed changes in wages, because changes in firm performance may reflect also market-level factors or changes in worker characteristics. Rather, to measure how sensitive wages are to firm-level—rather than market-level—conditions, it is necessary to study variation in firm conditions that are unrelated to the overall state of the labor market (“idiosyncratic”) and are unrelated to the characteristics of workers themselves (“exogenous”).

To estimate the degree to which wages adjust in response to firm conditions, we study the wage incidence of quasi-experimental, idiosyncratic shocks to product demand, which exogenously shift labor demand conditions across firms without impacting broader labor markets. To identify idiosyncratic product demand shocks, we study exposure of more than 4,000 Portuguese exporters to demand fluctuations in foreign markets.² We focus on demand shocks during the Great Recession, a period when global trade flows were sharply disrupted in

¹The term “internal labor markets” was introduced by Doeringer and Piore (1971), who highlighted three kinds of frictions in particular: 1) costly on the job training, 2) need for specific skills that are costly to find in thin labor markets, and 3) “customs,” i.e. norms, laws, and collective bargaining institutions that make it difficult for firms to replace incumbents with outsiders.

²This paper joins a growing literature combining rich firm-level export and import records with data on global trade flows to predict export demand and import supply shocks, for example Berman et al. (2015), Mayer et al. (2016), and Hummels et al. (2014). Our work is most related to the latter paper, which tests whether offshoring (input supply) shocks complement or substitute for incumbent workers. Their primary focus is substitutability/complementarity of low- and high-skill labor with off-shored inputs, though they simultaneously estimate import supply and export demand shocks as part of their analyses. They only report direct effects of (instrumented) log exports changes on wages, finding an elasticity of about 5% (if exports are a small fraction of sale, this presumably implies a larger sales-pass-through elasticity). This paper differs in its emphasis on identification shocks that are *idiosyncratic* to firms, and its focus on identifying the margins on which firms adjust their payroll, the magnitude of pass-through of shocks to total firm-level demand to wages, and the implications for wage determination.

unforeseeable ways, with significant variation in import behavior countries and across product markets. Because of unexpected changes in conditions across countries, exporters selling the same product but with pre-existing relationships with customers in different countries were exposed to differential changes in demand for their products. Crucially, this variation in exposure to different destinations within product markets generated differential shocks to demand for firms’ output—and to their demand for labor—without affecting the labor market as a whole.

Using detailed records on exports linked to balance sheets and matched employee-level data, we find that, after the onset of the recession, firms with better idiosyncratic demand shocks produce more output, expand their total payrolls, and pay their incumbent workers more. In contrast, we find no evidence that these shocks predict changes in production or payroll at *other* similar firms. To quantify the magnitude of these effects, we estimate a pass-through elasticity of wages with respect to sales that measures the causal change in wages per one-percent change in sales due solely to the demand shock. We show that, under weak assumptions, this elasticity is a lower-bound for the true elasticity between wages and the underlying change in firm-specific revenue productivity driving the changes in both sales and labor demand.³ We find that a shock that causes a ten percent increase (or decrease) in sales or value added results in a 1.5 percent increase (or decrease) in wages. Although sales at most firms declined during the recession, find that baseline nominal wage growth was nonetheless sufficiently inflationary that downward nominal rigidities did not appear to bind. We find that the pass-through effects are similar across all kinds of incumbent workers, though no effect is found for wages of new hires.

In our theoretical framework, we show that these results are difficult to reconcile with traditional models of competitive labor markets but are more consistent with a model featuring barriers that make it costly to fire incumbents and replace them with new hires. In line with such a model, we find that all employment adjustments occur on the hiring margin; there is no effect on differential retention. To further probe the plausibility of this explanation, we study heterogeneity in these pass-through effects across sectors where the extent of such barriers are likely to differ. In a wide class of models that microfound the barriers that give rise to internal labor markets, sectors with higher barriers to replacing incumbents will be marked by less labor market dynamism, lower turnover, and longer tenures. Accordingly, wage incidence should be most apparent in sectors with higher tenures and lower separation rates (higher “relationship-durability”).⁴ In sectors with lower separation rates, we find shocks that lead to a ten percent change in sales lead to an effect on wages exceeding three percent; by contrast, we find pass-through effects close to zero in low-durability sectors. These results follow the pattern of treatment effect heterogeneity one would expect to see in the model we outline, supporting

³By “revenue productivity” we refer to the combination of firm-specific demand and technical productivity (TFP) shifters, as in [Foster et al. \(2008\)](#).

⁴We classify firms according to the separation rate and typical tenure of permanent contract workers in the 5-digit industry. Examples of economic environments in which frictions generate higher relationship durability include models of firm-specific human capital following [Becker \(1962\)](#), such as [Jovanovic \(1979a\)](#) and [Lazear \(2009\)](#), models of heterogeneous match quality that require costly search, such as [Jovanovic \(1979b\)](#) and [Mortensen and Pissarides \(1994\)](#), and settings with institutional firing costs as in [Lazear \(1990\)](#) or convex costs to hiring [Acemoglu and Hawkins \(2014\)](#).

our explanation that wage incidence of firm shock reflects the extent of the barriers between internal and external labor markets.

These results provide novel evidence on the degree to which exogenous firm conditions can have substantial causal effects on wages. Although previous studies provide compelling evidence that product demand shocks to entire industries or regions have clear incidence on broad labor market segments (Abowd and Lemieux, 1993; Autor et al., 2014; Yagan, 2016), little is known about the causal incidence of firm-specific shocks onto the wages of individual employees. To identify the wage impact of firm-specific product-demand conditions it is necessary to isolate variation in firm performance arising from shocks that are *exogenous* to labor inputs provided by individual workers (uncorrelated with unobserved labor supply factors) and *idiosyncratic* relative to the external labor market (uncorrelated with unobserved market-level performance).⁵ Without a clear source of variation in firm performance, observed changes in output over time may arise in practice for many reasons and may have little relationship to changes in underlying product demand that in turn shifts internal labor demand. Our research design overcomes these obstacles. We find that the sensitivity of wage growth to firm conditions is large. Our elasticity estimates imply that moving an individual from a firm at the 25th percentile of three-year output growth to a firm that at the 75th percentile would have caused their wage growth to increase by 6.5 percent, roughly half of the interquartile range of the full wage growth distribution in our worker sample.

Our findings shed new light on the prior literature that associates observed differences in firm performance and observed changes in wages without a quasi-experimental shock to firm performance and help reconcile seemingly inconsistent findings. Previous work studying changes of firm performance over time to changes in employee wages have found a small but very robust relationship between changes in sales or production and changes in wages or earnings in longitudinal employer-employee matched datasets.⁶ However, Card et al. (2016b) note that while observational studies of longitudinal correlations between firm performance and employees' wages yield small elasticities (0.06 or less), the cross-sectional relationship between

⁵Some evidence does exist about the pass-through of firm-specific shocks to workers. One type of evidence comes from evidence on wage changes that occur after new innovations or patent awards. Van Reenen (1996) studies the effect of innovations on firm profits, and in turn the average wage bill of contemporaneous workers (one cannot observe whether or not the composition of workers changes in this case). Similarly, ongoing work by Kline et al. (2017) studies how patents and incumbent-worker earnings co-evolve after the award of a patent, exploring quasi-random variation in the timing of a patent. Both find significant effects of pass-through to workers' wages, though the later study finds the largest effects for those actually named on the patent. While the presence of wage increases tied to success in innovation and patenting contradicts fully competitive price-taking behavior by firms period-by-period, these successes are attributable to previous efforts of individual workers in prior periods (though whether or when the work will pay off is unknown in advance). By contrast, it is not clear from these studies whether a fully exogenous demand shock would have any incidence on workers. Additional evidence comes from the "pay-for-luck" studies by Bertrand and Mullainathan who found that executives without strong principals both receive higher pay when firms profit for reasons out of their own control, and also appear to pay workers more regardless of the work of the employees (Bertrand and Mullainathan, 2001, 1999). As before this work suggests that wages co-move with profits in the presence of contracting frictions, but does not directly study effects of internal *demand* for labor.

⁶See surveys by Manning (2011) and Card et al. (2016b). These include both studies that study changes in contemporaneous firm-level average wages inclusive of changing composition of workers, and within-spell wage changes for fixed sets of individual workers who stay at firms over time. For example, in related work Budd et al. (2005) study how average wages at firms change in response to both own-firm profits and the profits of foreign parents.

labor productivity and firm wages yields a relationship that is much larger. In particular, they find that firms with ten percent higher labor productivity have 1.2 percent higher firm pay premiums, reflected in firm fixed effects from a wage equation as in [Abowd et al. \(1999\)](#). Our results can partially reconcile these different findings. We find that OLS estimates of wage pass-through on the same sample yield estimates that are an order of magnitude smaller than the pass-through identified off exogenous demand shocks. We argue that this finding arises because observed changes in output are poor indicators of underlying product market conditions that determine labor demand, leading to attenuation of estimated pass-through of demand shocks.⁷ Thus, although market-level shocks and shocks to unobservable worker characteristics may bias observational estimates of rent sharing, in practice the first order concern for studies of rent-sharing that are not based on a well-defined firm shock may be that results are attenuated due to mismeasurement of labor demand.

The results relate to recent work that ties rising wage inequality to growing disparities across firms ([Furman and Orszag, 2015](#); [Card et al., 2013b](#); [Song et al., 2015](#); [Barth et al., 2016](#); [Alvarez et al., 2018](#)). In particular, our findings about the incidence of firm-specific trade shocks are consistent with recent work by [Helpman et al. \(2017\)](#), who show in an estimated structural model that increased dispersion in firm performance due to trade liberalization may directly result in increased wage dispersion within labor market segments. Moreover, our findings suggest that firms face substantial barriers to labor adjustment in response to shocks.⁸ Such labor-adjustment frictions may have important consequences for the aggregate performance of labor markets—for example, total job creation or the frequency and duration of unemployment spells—particularly during recessions or periods of high volatility.⁹

The paper is organized as follows: Section 2 presents the conceptual framework that illustrates how labor market imperfections give rise to wage incidence of firm-specific shocks and how the objects we estimate relate to underlying frictions in the labor market. Section 3 provides background information about the economic and institutional context, and it describes the data sources we use. Section 4 presents the empirical strategy of the paper and provides preliminary evidence for its validity. Section 5 presents reduced-form estimates of effects on

⁷This is similar the result found [Abowd and Lemieux \(1993\)](#) that observational changes in sales do not predict large wage changes, but sales predicted industry-level export shocks do predict large wage effects; [Card et al. \(2016b\)](#) note that simply instrumenting for output changes during a given interval with observed output changes over a slightly longer interval increases the estimated “rent-sharing” elasticity of wages to output from 3% to 6%.

⁸This supports the claim of [Oi \(1962\)](#) that labor may be better modeled as a “quasi-fixed” factor like plant capital than a fully flexible factor.

⁹[Mortensen and Pissarides \(1994\)](#) show how hiring and layoffs asymmetrically adjust to different kinds of shocks in the presence of heterogenous match quality that generates option value to maintaining matches. [Bertola and Cabellero \(1994\)](#) provide conditions under which labor adjustment costs lead cross-sectional inefficiencies due to labor hoarding, resulting in sluggish labor market performance. The papers by [Card et al. \(2013a\)](#) and [Braguinsky et al. \(2011\)](#) argue that the need to share the benefits of productivity improvements with workers via rent-sharing amount to a tax that discourages productive investment in Italy and Portugal, respectively. A large literature has studied how adjustment costs (particularly due to institutional firing costs) may have exacerbated slow job growth and high unemployment in Europe starting in the 1980s (“Eurosclerosis”), for example [Bentolila and Bertola \(1990\)](#) and [Blanchard and Summers \(1986\)](#). Recent work by [Schoefer \(2015\)](#) that distortions to hiring due to inability to adjust incumbent workers may be exacerbated in the presence of credit constraints.

firms—both the impact on production and revenues as well as the firm-level effects on labor demand and payroll adjustment. Section 6 presents the main results about pass-through effects onto incumbent workers’ wages. In Section 7, we test for heterogeneous pass-through across sectors with different degrees of relationship durability. In Section 8, we discuss how our findings relate to estimates from alternative sources of variation and to previous estimates of rent-sharing. Section 9 concludes.

2 Conceptual Framework

This section presents a theoretical framework relating the objects we estimate to underlying labor market features. In our framework, the incidence of firm-specific demand shocks on wages reflects the product of two factors: First, the share of wages due to noncompetitive rents (reflecting the costs of replacing incumbents with external hires); second, the extent to which shocks change the *relative* productivity advantage of incumbents over outside hires. To illustrate these two factors, we build a simple model of internal labor markets in which incumbent labor has a productive advantage over new hires from a perfectly competitive external market. In this specific setting, we show what conditions are necessary for firm-specific demand shocks to have incidence on incumbent wages, and we formally derive the pass-through elasticity we estimate. We then discuss how this interpretation extends to a broad class of models of wage determination in imperfectly competitive markets under more general assumptions.

2.1 A Stylized Model of Wage Incidence in Internal Labor Markets

2.1.1 Set Up

We begin with a stylized setup in which firms employ workers that are costly to replace after their hiring. We consider a set of firms requiring the same skill type, such that all employers and workers considered participate in the same labor market segment and all have access to a common competitive market.¹⁰

Timing. The model takes place during a single period. The sequence of events during the period as follows: First, firms inherit demand and productivity attributes, as well as a stock of incumbent employees, and then realize a shock to their demand level. Second, firms negotiate wages with incumbent workers. Third, firms decide on the number of external hires they will employ. Finally, production occurs and wages are paid at the end of the period.

Firms. Firms j are small relative to the labor market and begin the period with a mass L^{inc} of incumbent workers. They produce revenue R_j using L^{inc} incumbent workers and a mass L^{out} of hires from the outside market:

$$R_j = P_j \times \bar{P} \times A_j \times f(L^{inc}, L^{out}) \quad (1)$$

¹⁰For example, all firms may be in a single industry that uses an industry-specific skill. The important simplification is that all firms utilize the same skill type; it may be the case that subsets of firms operate in different product markets.

In this set-up, revenue is the product of $f(L^{inc}, L^{out})$ (assumed to be increasing and concave in its arguments) the total factor productivity (TFP) A_j of the firm, the market-level price index $\bar{P}$, and a firm-specific price shifter P_j . In this set-up, firm-specific demand shocks and TFP shocks enter the firms' problem symmetrically through a common *revenue productivity* term $P_j \times A_j$, as in [Foster et al. \(2008\)](#), so proportional shocks to either have identical effects. For the following exposition, we consider behavior when A_j and $\bar{P}$ are fixed and suppress these variables.

A key feature of (1) is that incumbent labor and outside hires are not assumed to be perfect substitutes. As in [Doeringer and Piore \(1971\)](#), firms may face barriers to replacing internal incumbent workers with equivalent external hires, and workers may likewise face barriers finding other equally-suitable jobs.¹¹ In particular, firms may have to forgo valuable production time in order to substitute new hires for incumbents if doing so requires time recruiting and training replacement workers with specialized skills, negotiating with unions with strike power, or engaged in legal proceedings triggered by labor laws. As a result, external hires are not perfect substitutes for incumbent workers, who have a strictly higher net marginal revenue product (*inclusive* of replacement and retraining costs): $f_1(L^{inc}, L^{out}) > f_2(L^{inc}, L^{out}) \forall L^{inc}, L^{out}$.¹² In Appendix A, I illustrate how this production function can arise from a standard production function using homogenous labor with firing and hiring costs.

Labor Markets and Wage Determination. Firms hire L^{out} homogenous new workers from a competitive external labor market, with going wage $\bar{w}$.¹³ The competitive going wage $\bar{w}$ is determined in equilibrium by the total labor supply (assumed to be fixed) and the total derived demand for hires, which reflects the prevailing market-level product demand level $\bar{P}$ (thus $\bar{w} \equiv \bar{w}(\bar{P})$). Incumbent workers at firms always have the option of entering the external labor market and earning $\bar{w}$; the outside wage $\bar{w}$ is taken as parametric by both individual firms and workers, who are small relative to the market.

By contrast, the stock of incumbents L^{inc} is exogenously determined. We make two important assumptions about incumbent workers: First, L^{inc} is sufficiently low relative to $P_j \times \bar{P} \times A_j$ that firms always want to make positive hires $L^{out} > 0$ even when they retain all incumbents L^{inc} .¹⁴ A direct result of this assumption is that for some firm-specific wage $w_j > \bar{w}$, the firm is strictly better off by retaining all employees than by replacing incumbent workers with external hires. Second, we assume that the set of non-retiring incumbents bargain over their wages (but not the level of firm hiring) collectively as a unit; thus, all incumbents are either retained at a settlement

¹¹Following [Doeringer and Piore \(1971\)](#), theoretical work has highlighted three primary types of worker replacement costs that create barriers between internal and external labor markets: 1) costs of rapidly replacing firm-specific know-how acquired through experience [Becker \(1962\)](#); 2) costs of having to find replacements for good incumbent matches [Jovanovic \(1979b\)](#); or 3) direct institutional barriers to firing due to laws or unions ([Lazear, 1990](#); [Bertola and Cabellero, 1994](#)).

¹²Studying employer responses to unexpected deaths of incumbent employees, [Jaeger \(2015\)](#) provides evidence that internal and external workers are not perfect substitutes.

¹³In this one-period model, we assume hires can be costlessly replaced during the recruiting process, so there is no wage incidence on hires. However, this is not an important assumption for the results that follow. In more general models, hires may become costly to replace immediately upon contact, as in models with costly search.

¹⁴In a fully dynamic model, this might reflect an exogenous retirement rate that leads firms to always make positive hires in steady state.

wage w_j or leave as a bloc if the bargaining process falls apart.¹⁵

We assume wages are determined as the generalized Nash bargaining solution, where the incumbents and the firm have bargaining weights γ and $(1 - \gamma)$ respectively. Specifically, the settlement wage for incumbents at firm j , w_j solves

$$w_j = \underset{\omega}{\operatorname{argmax}} (\text{Firm Surplus}(\omega))^{(1-\gamma)} \times (\text{Combined Incumbent Surplus}(\omega))^\gamma \quad (2)$$

where the respective surpluses given wage w_j relative to the outside options will be derived below.

2.1.2 Firm Behavior and Wage Determination

If negotiations with incumbents result in agreement on a settlement wage w_j , firm profits (given optimal hiring) are¹⁶:

$$\Pi_j^1 = \underbrace{\max_{L^{out}} \{P_j \times f(L^{inc}, L^{out}) - \bar{w}L^{out}\}}_{V(L^{inc}, P_j)} - w_j L^{inc} \quad (3)$$

However, if negotiations fall apart, firm profits are:

$$\Pi_j^0 = \underbrace{\max_{L^{out}} \{P_j \times f(0, L^{out}) - \bar{w}L^{out}\}}_{V(0, P_j)} \quad (4)$$

While the choice of L^{out} does not depend on the settlement wage w_j , it does depend on whether or not the firm has retained the incumbent workers.¹⁷ Hence, one can write firm profits in terms of the net revenues before incumbent wages are paid $V^1(L^{inc}, P_j)$ or $V^0(L^{inc}, P_j)$, which depend on the demand level P_j but not the settlement wage w_j , and the payment to incumbents $w_j L^{inc}$ that occurs if an agreement is reached.

The firms' surplus from retaining the incumbents can be written in terms of the two indirect profit levels combined with the negotiated incumbent wage bill:

$$\text{Firm Surplus}(w_j) = V(L^{inc}, P_j) - V(0, P_j) - w_j L^{inc} \quad (5)$$

The incumbents' surplus is the excess of the wage bill over what they could earn in the outside market

$$\text{Combined Incumbent Surplus}(w_j) = w_j L^{inc} - \bar{w} L^{inc} \quad (6)$$

Given the bargaining weights γ and $(1 - \gamma)$ for workers and the firm, respectively, the settlement

¹⁵This is similar to wage determination in the "Insider-Outsider" model of [Lindbeck and Snower \(2001\)](#). In this simplified variant, we assume workers negotiate only over their own wages, and not the hiring level itself (as modeled by Lindbeck and Snower). Hiring decisions are made unilaterally anticipating the resulting wage bargain with incumbents, similar to "right-to-manage" union models following [Dunlop \(1950\)](#).

¹⁶We have suppressed $\bar{P} \times A_j$, which are assumed for now to be fixed multipliers of P_j .

¹⁷The choice of L^{out} is the same for any level of w_j , since hires never earn incumbent wages in the one-period model, and since incumbents do not directly negotiate hiring levels jointly with the wage.

wage that solves (2) is thus:

$$\begin{aligned}
w_j &= \bar{w}(\bar{P}) + \gamma \times \frac{1}{L^{inc}} \times [V(L^{inc}, P_j) - V(0, P_j) - \bar{w}L^{inc}] \\
&= \underbrace{\bar{w}(\bar{P})}_{\text{Outside Option}} + \underbrace{\rho(P_j)}_{\text{Non-Competitive Quasi-Rent}}
\end{aligned} \tag{7}$$

where $\rho(P_i) = \gamma \times \frac{1}{L^{inc}} \times (V(L^{inc}, P_j) - V(0, P_j) - \bar{w}L^{inc})$ is the part of the wage comprised of a noncompetitive quasi-rent. This latter component is the share each worker captures of the total surplus from maintaining the employment relationship, relative to each parties outside options. Intuitively, the worker earns the competitive wage achievable on the market, plus a premium—a “rent-sharing” term—reflecting excess internal demand for incumbent labor, due to frictions in the labor market.

Proposition 1 summarizes how idiosyncratic shocks impact the quasi-rent sharing term:

Proposition 1. Under the assumptions of the model, when $f_1(L^{inc}, L^{out}) > f_2(L^{inc}, L^{out})$, the quasi-rent-sharing term $\rho(P_i)$ always satisfies $\rho(P_i) > 0$ and $\rho'(P_i) > 0$.

Proof. See Appendix A. □

The first result (positive rents) follows directly from the productive advantage of incumbents over external hires. The second result (rents increase in firm demand) occurs because the *revenue* productivity advantage of incumbents is augmented by demand P_j in our model. If incumbents had a revenue productivity advantage that was positive but invariant to demand conditions—for example, if replacing workers required firms to incur a constant currency cost—this second result would not hold. By contrast, our formulation of (1) formalizes the intuition that the time spent replacing an incumbent with an outside hire is costlier when demand conditions are better.

As long as the conditions in Proposition 1 hold, firm-specific shocks affect incumbent wages through the rent-sharing term $\rho(P_i)$.¹⁸ By contrast, first term $\bar{w}(\bar{P})$ depends only on market-level demand—it can therefore change in response to common demand shocks to $\bar{P}$ but not idiosyncratic demand shocks to P_j .

For comparison, the perfect competition benchmark is reflected in the case where incumbents and external labor are perfect substitutes, so that $f(L^{inc}, L^{out}) = f(L^{inc} + L^{out}, 0) = f(0, L^{inc} + L^{out})$ and therefore $f_1(L^{inc}, L^{out}) = f_2(L^{inc}, L^{out})$. When incumbents can perfectly substitute for external hires, the total surplus shrinks to zero, and, thus, the settlement wage per incumbent is $\bar{w}$. Intuitively, in a perfectly competitive market, firms can always find replacement work at wage $\bar{w}$ —thus, they never pay more. In this case, when firms experience demand shocks of any kind, payroll adjustment only occurs on the hiring margin; wages only reflect market-level demand (and hence *common* shocks to $\bar{P}$) but are invariant firm-specific labor demand.

¹⁸Under more general assumptions, it is *not* necessary that $\rho'(P_i) > 0$ is the case whenever $\rho > 0$. If $\rho'(P_i) = 0$, then wage will be zero despite the presence of positive rent-sharing.

2.2 The Pass-Through Elasticity

In the model above, the effects of idiosyncratic demand shocks on wages, holding $\bar{P}$ and $\bar{w}$ are characterized by the following pass-through elasticity:

$$\epsilon^{w, P_j | \bar{P}} \equiv \frac{\partial \ln w_j}{\partial \ln P_j} = \underbrace{\frac{\rho(P_j)}{w_j}}_{\text{Rent Share in Wage}} \times \underbrace{\frac{d \ln \rho(P_j)}{d \ln P_j}}_{\text{Change in Inside vs Outside Value}} \quad (8)$$

This elasticity is the product of two terms. The first term is the share of the wage that is comprised of the rents arising due to costs of replacing incumbent workers. Pass-through effects will therefore be larger when noncompetitive rents comprise a larger portion of wages, reflecting larger deviations from the benchmark of perfect competition. The second term is the elasticity with which the per-worker surplus (i.e., the relative demand for incumbents relative to outsiders) changes with respect to a productivity shock. Because of the presence of this term, no pass-through effect exists, even in cases where wages reflect positive rents, if the quantity of rents is invariant to demand shocks. Intuitively, this second term reflects how shocks change the *relative* productivity advantage of incumbents over outside hires due to imperfect substitutability.

This pass-through elasticity is not specific to the particular model derived above, in which incumbent labor is supplied inelastically and external hires are paid a competitive wage. Rather, a wide class of models of imperfect labor market competition yield wage equations of the form $w_j = \bar{w} + \rho(P_j)$. First of all, this type of wage equation exists in most bargaining models where workers cannot be costlessly replaced in the short run, regardless of the specific bargaining protocol. These include the single-worker firms in the Diamond-Mortensen-Pissarides model (Pissarides, 2000), sophisticated models of multilateral non-cooperative bargaining between individuals and employers with hiring costs (Stole and Zwiebel, 1996; Acemoglu and Hawkins, 2014), and models of union bargaining (e.g. Brown and Ashenfelter (1986)). In the textbook monopsony model, firms also post wages and face a firm-specific upwards-sloping supply curve due to heterogenous worker tastes, as in Manning (2011). Similar to our model, the unobserved heterogeneity in match quality that generates the upwards-sloping supply curve implies that incumbents are not perfectly replaceable by outside hires.¹⁹ A similar wage equation occurs in efficiency wage models in which incumbent effort has higher returns than external hiring, and firms face incumbent effort curves that are upwards-sloping in wages (Katz, 1986).²⁰ In this latter

¹⁹Unlike our model, the baseline wage-posting model (where all employees receive the same posted wage) has the feature that increases in wages apply to all workers, including hires, and are directly related to increases in employment levels. However, one can write a similar version of our model where new hires are supplied with infinite elasticity, but incumbent *retention* is an upwards-sloping function of the incumbent wage. In this case, higher incumbent wages implement higher incumbent retention rates, and need not be reflected in the wages of hires. We find some evidence in support of this model, discussed below.

²⁰As noted below, $\rho'(P_j) > 0$ only occurs if the marginal product of incumbent effort is higher than effort of new hires. Relatedly, models of non-contractable effort following Holmstrom (1979) can give rise to similar wage equations when it is difficult for firms to determine whether changes in firm performance are due to product demand shocks or

case, increases in unobserved labor inputs would be empirically indistinguishable from increases in the wage for fixed quantities of efforts and would therefore be interpreted as rent-sharing in observed wage data.

In each of these models, the pass-through elasticity is a reflection of the same two properties of the labor market: the *level* of frictions reflected in the share of wage that comes from rents $\frac{\rho}{\bar{w}}$ and the imperfect substitutability of incumbents and external hires reflected in the degree to which the surplus per incumbent changes with the demand level $\frac{d \ln \rho}{d \ln P_j}$. In all cases, the pass-through elasticity tends towards zero as the underlying friction diminishes and the labor market tends towards perfect competition. In addition, the elasticity tends to zero when incumbent labor and marginal external hires are perfect substitutes, even when there is positive rent-sharing ($\rho > 0$).²¹

The elasticity in (8) is highly similar to the “rent-sharing elasticity” defined by Card et al. (2016b), which is simply the second term $\frac{\rho}{\bar{w}_j}$. The difference is that the latter quantity is the elasticity with respect to the *surplus* ρ , assuming it is known; by contrast, we do not assume the structural surplus is observable, so the expression in (8) incorporates the “first-stage” change in surplus. Importantly, positive pass-through therefore requires that the per-worker rent ρ changes due to the shock. Even if $\rho > 0$, if ρ is invariant to shocks, then pass-through elasticities cannot be informative about ρ .

Importantly, this framework highlights that $\epsilon^{w, P_j} | \bar{P}$ can only be identified from shocks to firms that are *exogenous* to labor supply factors and *idiosyncratic* to the firm. While (8) yields the interpretation of pass-through of idiosyncratic demand shocks, it does not characterize wage changes due to common shocks to $\bar{P}$. Common shocks have an *additional* effect on the wage through the outside option wage $\bar{w}$, which workers earn regardless of firm conditions. Elasticities identified off of market-level variation may reflect some arbitrary combination of firm-level and market-level pass-through; one could only disentangle the two effects under strong additional assumptions. In addition, $\epsilon^{w, P_j} | \bar{P}$ is not identified in the data if increases in output and wages occur due to labor-supply-side shocks (e.g., effort per worker), rather than firm-side factors (P_j). The elasticity in (8), only includes wage changes due to of firm-side shocks.²²

While (8) defines the elasticity of wages with respect to a firm-specific shock to the demand

employee performance. In the Holmstrom (1979) formulation of the incentive contract, firms will pay workers a base wage (reflecting market competition) plus some amount that varies with firm performance.

²¹The cases where rents are positive ($\rho(P_j) > 0$) but invariant to shocks ($\frac{d \ln \rho(P_j)}{d \ln P_j} = 0$) are those where all employees earn rents, but new hires and incumbents are perfectly substitutable. For example, Abowd and Lemiue (1993) show that this can occur in the strongly-efficient union bargains model of Brown and Ashenfelter (1986) where unions maximize total member payroll *including* workers in the union but outside the firm. Since the union views payroll of incumbents and payroll of union members who are firm outsiders as perfect substitutes, union bargains may accommodate number of workers while keeping the *per – worker* rents exactly constant (this occurs when the production function is iso-elastic / Cobb-Douglas in labor). A similar result occurs in the classic efficiency wage model of Solow (1979), in which firms maximize total effort eL , where e is effort per worker and is increasing in the firm’s wage level. In the model where firms maximize eL —so that incumbent effort and effort of hires are perfect substitutes—wages are a constant markup over outside wages that is invariant to productivity shocks. However, if incumbent effort could not immediately be replaced by new hires, this invariance to shocks would no longer hold.

²²Increased labor supply by incumbents need not change the firms’ output, however, as firms may adjust its labor stock on other margins to keep the total number of efficiency units constant. This would occur in the perfectly competitive benchmark, for example.

level P_j (or, equivalently, TFP A_j), P_j is not directly observed in practice. In our analysis, we study changes in observed output, which can be used as a proxy for revenue productivity. We quantify how much firm-specific demand shock yielding a 1 percent change in observed output passes through to wages. Using output as a proxy for P_j will generally produce a biased estimate of the true $\epsilon^{w,P_j}|^{\bar{P}}$; however, the model above allows us to sign this bias using weak assumptions. Whenever firms weakly increase (or decrease) output in response to a positive (or negative) shock to P_j in *excess* of what would occur if they kept inputs constant, $\frac{d \ln R_j}{d \ln P_j} \geq 1$, so that percent changes in total output weakly *overstate* changes in demand.²³ For example, in the setup above, firms always weakly increase (or decrease) the target employment level L^* in response to positive (or negative) shocks. Thus, $\frac{d \ln R_j}{d \ln P_j} = 1 + f'(L) \times \frac{dL}{dP_j} \geq 1$. Although we believe benchmarking wage changes to the output change caused by the same shock is useful in its own right, our estimates also can be interpreted as a lower bound for $\epsilon^{w,P_j}|^{\bar{P}}$.

3 Setting and Data

3.1 Portuguese Exporters in the Great Recession

Our analysis exploits variation in export demand experienced by Portuguese exporters during the Great Recession of 2008 and 2009. From 2007 to 2009, an unforeseen housing bust in the United States and the resulting financial crisis triggered large recessions around the world, precipitating a sharp import demand drop in many countries in 2009 and 2010 (Eaton et al., 2016). In Portugal, however, there was no major domestic housing or financial crisis that directly affected firms; nonetheless, they were highly exposed to global fluctuations in trade demand. Portugal is a small, open economy: while Portugal's exports comprised only 11 percent of its GDP in 2007, exports amounted to over 31 percent of output (World Bank, 2016). Figure 1 highlights how the recession that occurred in Portugal during this episode was marked by a dramatic decline in total exports that mirrored global trends; indeed, exports did not return to pre-recession levels until 2011. During this initial recession, unemployment grew moderately, at a pace also comparable to other advanced European economies. By contrast, after the economy had entered recovery at the end of 2010, Portugal experienced a sovereign debt crisis during 2011 that resulted in a second, distinct recession episode, which then culminated in 2012 with dramatic increases in unemployment. To avoid concerns about confounds arising from this latter, domestic crisis, our study focuses on demand variation that occurred during the first recession episode through 2010.

²³Changes in *per – worker* output may have *no* similarly well-defined relationship to the underlying change in P_j . In many cases, changes in $\ln \frac{R_j}{L_j}$ may significantly *overstate* changes in underlying revenue productivity—indeed, in standard models $\ln \frac{R_j}{L_j}$ remains *constant* in response to any change in revenue productivity. For example, this occurs when production is isoelastic (i.e. Cobb-Douglas), so that the marginal product of labor is always in a constant ratio to *average* output per worker, and firms efficiently adjust employment quantities in response to idiosyncratic shocks to keep the marginal product of labor equal to the (constant) outside wage. For any change in underlying productivity, average product (output per worker) remains constant.

While Portugal provides an ideal setting to identify idiosyncratic demand shocks, key features of labor market institutions should be taken into account when considering the external validity of our results. Portuguese labor markets are substantially more rigid than those in the United States; prior to the Great Recession, [Blanchard and Portugal \(2001\)](#) observed that while steady-state unemployment rates were actually quite similar in the United States and Portugal (in contrast to other Southern European nations), this similarity masked stark differences in dynamism. Entry into unemployment was much lower in Portugal, but durations were much higher. One potential explanation for the low degree of turnover in Portugal is the presence of very strong job protections. Similar to systems in several other European countries (e.g., France), most employees work under permanent contracts that can only be terminated if the firm has legally defensible cause. Popular accounts suggest layoffs are thus extremely rare, and the lack of an option to downsize results in higher hesitation of employers to create new jobs, particularly in the presence of uncertainty about future economic conditions.

In recent decades, firms have been granted limited scope to hire workers under fixed-term contracts, wherein after a set duration employees must either be released or promoted to a permanent contract²⁴. Prior to the Great Recession, fixed-term contracts accounted for about 15 percent of total private-sector employment. While these contracts offer firms some degree of flexibility, they are nonetheless still quite rigid compared to typical work arrangements in the United States because firms may not dismiss workers without cause until the full duration has elapsed. Instead, these contracts offer firms an opportunity to learn about the quality of a match before committing to a permanent contract; accordingly, they fully account for 50 percent of all hires in Portugal.

Collective bargaining over wages is widespread in Portugal, with over 90 percent of private sector employees covered by bargained contracts between 2010 and 2012. Collective bargaining institutions are distinctive: although only 11 percent of private sectors are unionized, under Portuguese law, any wage floor negotiated for a specific job title automatically extends to all workers with the same job title.²⁵ Firms can freely raise wages above these floors but have very limited scope to reduce nominal wages.

3.2 Data Sources and Sample Selection

To study how firm-level product market conditions pass through to wages, we use anonymized firm identifiers to link detailed administrative records on exports, balance sheets, and profit-loss statements to matched employer-employee personnel records. The export data is derived from administrative customs records for exports outside the European Union and from mandatory

²⁴Fixed term contracts may be renewed up to two times under certain circumstances, but were not indefinitely renewable during the study period

²⁵Unionization and contract coverage figures are from [Addison et al. \(2017\)](#). [Martins \(2014\)](#) provides a detailed discussion the institutions that result in extension of bargained contracts—resulting in over 30,000 occupation-specific wage floors—and studies the impacts on the wage distribution.

reporting on all intra-EU shipments in excess of a certain threshold.²⁶ These data, which cover the period from 2005 to 2013, report export and import volumes at the firm by destination country by six digit product (HS-6) level. To relate firm behavior to demand condition in foreign markets, we use annual data on trade flows by six-digit product and country pair publicly available in the BACI database.²⁷

We link the export records to profit and loss statements and to balance sheets for the universe of firms operating in Portugal contained in the *IES* database, with coverage for the years 2005–2013. This database includes information on the sales, input costs, factor payments, profits, assets, and liabilities of all Portuguese firms on an annual basis. We use items in this dataset to calculate the value added of the firm, defined as the the output of the firm net of purchases of intermediate goods in services²⁸. While gross sales are a useful measure of output as they are directly and reliably observed, value added is a more direct measure of the level of production in the firm.

We then use the same anonymized identifiers to link firms to the the Quadros de Pessoal (QP), a matched employer-employee dataset produced by the Ministry of Employment based on a census of all private-sector employers during October of each year. Employers with at least one paid employee must report each employee’s baseline monthly contract wages, fringe payments, hours worked (regular and overtime), gender, detailed occupation, tenure, age, contract type, and education level of every worker actively employed during the reference month²⁹. We define workers as *incumbents* at firm j if they worked full-time at firm j in 2007 and if this was their primary source of reported earnings. Using longitudinal employee data enables the study of effects on the incumbent cohort over time, even as workers may move to other firms in subsequent years. In the primary analysis, we study effects on “attached” incumbents, where workers are deemed to be attached to the firm if they are observed working full-time at firm j in 2005, 2006, and 2007; however, results for less-attached incumbents are presented as well. Because reported quantities pertain to a single reference month during the year, the primary compensation outcome is the base monthly contract wage, excluding fringe payments and overtime which may be inconsistent over the year, as this is most likely to stay constant throughout. We also study effects on hourly wages, calculated as the base wage divided by the regular hours specified in the employment contract.

Our analysis sample consists of all pre-period exporters subject to two restrictions. First, we

²⁶110,000 Euros in 2007.

²⁷ [Gauiler and Zignago \(2010\)](#) overview the underlying source data and discuss how the harmonized trade flow database was constructed.

²⁸In practice, calculation of value added during a specific period requires one to subtract the intermediates used in production of goods sold during that period regardless of when the intermediates are purchased—this requires an adjustment for inventory stocks. Since value added is defined as the total output of labor and capital factors in the firm during the year, an accounting identity requires that the total value added of the firm equal the total factor payments (all costs of employed labor plus payments to capital including net profits and taxes on profits). Since the data on factor payments are more reliable than the data on inventories in the data, we follow [Card et al. \(2016b\)](#) and define value added based on the latter concept, so that $VA = \text{total labor costs} + \text{gross earnings before netting out interest, taxes, depreciation, and amortizations}$.

²⁹Government workers, workers temporarily out of the labor force, or working as an independent contractor are excluded from these data

restrict to firms that both export and employ more than one worker in 2005, 2006, and 2007—these firms are employers likely to export in future years. Second, we limit our primary focus to small- and medium-sized firms—those with average 2005–2007 employment of at least one and no greater than 100 employees—as these firms are most plausibly exposed to *idiosyncratic* demand shocks based on skewed exposure to specific destinations. Larger firms are observably more diversified across customer markets and are less exposed to idiosyncratic demand risk. Although the 1,000 largest firms excluded from the analysis account for a majority of sales and employment among exporters, the 4,178 firms in the analysis firms are most representative of typical private sector employers in Portugal.³⁰ Appendix Figure A.1 shows that, in tabulations of the 2007 data on all firms, the firms in the sample are actually much larger than most firms in Portugal; however, when tabulating the distribution of *workers*, the firms in the sample are central in the distribution of firm sizes. The sample includes all sectors that export goods; while the majority of firms are manufacturers, the sample also includes firms in resource-extraction industries, wholesale and retail, and select service industries that produce intellectual property (such as book or software publishing).

Table I presents summary statistics of the pre-period export behavior of sample firms. On average, exports reported in the shipment-level data (subject to reporting thresholds) account for 34 percent of sales reported on the firm’s profit/loss statement, though the median is smaller and approximately 20 percent. Firms in this sample are not highly diversified across destinations—the median firm only exports to three countries. Similarly, firms are not highly diversified across products; the median firm only exports two major products and four detailed varieties, whereas, by contrast, the larger export firms are notably more diversified and sell more products to more destinations.

Table II presents summary statistics of the workers in the sample in 2007, as compared to the full population of full-time workers in the QP. The distribution of earnings, hours, and worker demographics in the sample closely matches the distribution workers in the broader private sector. Workers at exporting firms have somewhat longer tenures—the median worker in the sample has been employed at the same firm for approximately three years more than the typical worker in Portugal. In both the sample and the broader population, the standard deviation of log wages is 0.5.³¹ By contrast, there is almost no variation in the normal working hours among employees in the data.

4 Empirical Strategy for Identifying Idiosyncratic Shocks to Demand

To estimate the pass-through of idiosyncratic demand shocks to the wages of the incumbent cohort of workers present at the firm before the recession, we implement a dynamic difference-in-differences design, in which we compare differential changes in incumbent wages

³⁰Exporting firms as a group tend to over-represent larger firms.

³¹the statutory minimum monthly full-time wage of 403 Euros (2007) is not binding for most workers in the sample.

across exporters that that experienced different idiosyncratic demand shocks during the Great Recession.³² In this section, we begin by describing how we use detailed export records to construct a firm-level predictor of export demand shocks during the recession. Second, we discuss how we decompose this initial shock into a component that is common to many producers and a component that we argue is idiosyncratic to each firm. Third, we present the estimating equations we use to study reduced-form effects and pass-through to workers. Finally, we summarize the identification assumptions that must hold in order to identify the objects of interest and offer evidence in support of these identifying assumptions.

4.1 Predicting Export Demand

We isolate exogenous, idiosyncratic changes in demand for firms' products by studying quasi-experimental shocks to export demand during the Great Recession. To construct firm-level shocks to demand that do not mechanically reflect changes in firm productivity, we construct a firm-level "shift-share" predictor based on exposure to foreign markets.³³ We measure the exposure of exporting firms to different destination markets prior to the Great Recession during the years 2005–2007, then, based on this fixed exposure, we apportion to firms the changes in import demand in each country-by-product market from before the recession (2006–2007) to the trough of the decline in trade (2009–2010) Because we measure for destination-market demand using observed changes in imports from everywhere in the world *excluding* Portugal, differences in demand are not mechanically reflective of changes in Portuguese export productivity to those markets.³⁴

Formally, we measure the exposure weight of a firm j to the market for six-digit (HS6) product p in country c as their share $s_{j,pc}$ of exports of product p to country c in the total 2005–2007 exports by the firm across all products (in set P) and countries (in set C):

$$s_{j,pc} = \frac{Exports_{j,pc}^{2005-2007}}{\sum_{p \in P, c \in C} Exports_{j,pc}^{2005-2007}} \quad (9)$$

The denominator is simply the total pre-period exports of firm j .

To construct an index of import demand changes in each destination market, we calculate the proportional change in imports of product p by country c from all countries, excluding Portugal between the two years before the global recession (2006 and 2007) and the two trough years of the

³²we study effects on the wages of a fixed cohort to ensure wage effect at the firm are not simply a product of changes in composition of the types of workers employed. we consider effects both unconditional and conditional on workers' remaining at the initial firm, to account for the possibility of rent losses due to switching jobs.

³³Our shock is may be considered a firm-level Bartik shock, using firm-level exposure weights to country-by-product markets.

³⁴There may be a second-order effect of productivity shocks to Portuguese exports arising from crowd-out of competitors. In this case, an observed decline in non-Portuguese imports could reflect growing market share of Portuguese firms. To the extent that growth in observed non-Portuguese imports in a market actually predict growth—not declines—in Portuguese exports to that market, this channel is small relative to the first-order demand prediction.

global decline in trade (2009 and 2010). Since it is possible that some countries stopped importing some products altogether during this period, we approximate the percentage change using the symmetric growth rate (or “arc-elasticity”) concept commonly used in literature on firm dynamics [Davis et al. \(1996\)](#). Denoting the average 2006 and 2007 level of non-Portuguese imports (NPI) of product p to country c as NPI_{pc}^{pre} and the corresponding average level during 2009 and 2010 as NPI_{pc}^{post} we calculate the change in imports Δ_{pc} as

$$\Delta_{pc} = \frac{NPI_{pc}^{post} - NPI_{pc}^{pre}}{\frac{1}{2}(NPI_{pc}^{post} + NPI_{pc}^{pre})} \quad (10)$$

The baseline predicted change in export demand for firm j , Δ_j , is calculated as the average change in each destination (country by product) market, weighted by the pre-period exposure of firm j to that market:

$$\Delta_j = \sum_{p \in P, c \in C} s_{j,pc} \Delta_{pc} \quad (11)$$

If valid, the shock Δ_j will predict changes in firm j ’s exports due strictly to changes in product demand.

4.2 Isolating Idiosyncratic Variation in Demand

The baseline shock Δ_j reflects two types of variation; one kind that is likely to impact entire markets, and another that is plausible idiosyncratic to individual firms. First, Δ_j reflects changes in demand for product p that occur in all countries and therefore may affect all producers of p in Portugal. Second, Δ_j also reflects differential demand adjustments across countries c for the *same* product p . Formally, let $s_{p,j} = \sum_{c \in C} s_{pc,j}$ be the total share of pre-period exports by firm j that are of product p (regardless of the destination), and let Δ_p denote the common decline in demand for product p across all global markets. Then Δ_j can be decomposed as:

$$\Delta_j = \underbrace{\sum_p s_{j,p} \Delta_p}_{\text{Common Cross-Product Change}} + \underbrace{\sum_{p,c} s_{j,pc} (\Delta_{pc} - \Delta_p)}_{\text{Differential Dest. Change Within Product}} \quad (12)$$

So long as firms use the same workers to produce p regardless of the destination, the latter component reflects idiosyncratic demand variation.³⁵

Accordingly, our approach is to decompose the predictor Δ_j , tabulated in (11), into two components: the part reflecting product-level demand changes common to many firms and the residual variation in Δ_j orthogonal to that common component. To do this, we attempt to directly

³⁵It does not need to be the case that no two firms in the labor market have the same shock, but rather that shocks to firm j do not predict product demand changes at enough other firms in the same labor market to change the wage. This might be violated if firms producing the same product p sort into distinct labor markets corresponding to the location of their customers—if firms exporting product p to different countries c draw from different labor markets corresponding to their destination *and* product, then there is no idiosyncratic variation in Δ_j .

measure changes in product-level import demand Δ_p common to all countries based on the bilateral trade flow data.³⁶ We proxy for the common demand shock to product p using the observed change in imports of product p averaged across all countries. In the baseline analysis, we use a simple unweighted average taken across all countries : $\hat{\Delta}_p = \frac{1}{\#c} \sum_c \Delta_{pc}$.³⁷ Given proxies for the product-level demand change common to all countries, we can then construct the predicted change in export demand based solely on the products exported by the firm ($s_{j,p}$) and the common change in demand for those products ($\hat{\Delta}_p$). We denote this initial “product-only” demand predictor for firm j as $C_j^0 = \sum_p s_{j,p} \hat{\Delta}_p$.

Next, we isolate the variation in Δ_j that is orthogonal to the product-only demand predictor C_j^0 and higher order terms. If one were to simply subtract C_j^0 from Δ_j , as in (12), the resulting residual may still be correlated with the product-only predictor in samples where firms do not export to all destinations.³⁸ Hence, to obtain the idiosyncratic component we residualize Δ_j on a quartic polynomial in C_j^0 . In doing this, we attempt to remove all variation in Δ_j that can be predicted based solely on what products j makes and the common demand change for those products.

The result is an empirical decomposition of Δ_j into two orthogonal components: first, a common component C_j predicted by C_j^0 and higher-order terms and, second, the residual *idiosyncratic* component, which we use as our main shock and label S_j . The idiosyncratic S_j only reflects differences in demand across countries for the same product that are *uncorrelated* with product-level demand changes.

Figure 2 plots the distribution of both the baseline demand predictor Δ_j and the idiosyncratic shock component S_j . In practice, the distributions of the two variables are very similar, with S_j accounting for 87 percent of the variation in Δ_j and the common component C_j accounting for only 13 percent. While the mean of the idiosyncratic shock is zero by construction, the zero mean of the baseline demand predictor is not mechanical but rather reflects the true distribution of export demand flextions faced by firms. Although the typical predicted demand change is only slightly negative, it is useful to compare these changes to typical pre-period demand changes over the same horizon. Figure 2 therefore also plots the distribution of the demand predictors calculated for each firm j holding the exposure weights fixed, but using the import demand change from 2003 and 2004 to 2006 and 2007 at firm j 's destinations, reflecting the same time duration as the baseline shock. In comparison to the average 28 percent pre-period growth of demand in firms' overseas markets, the recession shocks look markedly more grim. Accordingly, it is most reasonable to

³⁶An alternative approach would be to match firms with identical product compositions and study changes in demand predicted by cross-country differences alone. Matching firms with identical product portfolios is simply a method for identifying the common variation $\sum_p s_{j,p} \Delta_p$ without measuring each Δ_p directly. Since matched firms have identical $s_{j,p}$, then $\sum_p s_{j,p} \Delta_p$ would be absorbed by a match-level fixed effect—each firm in the pair acts as a counterfactual for the other. In practice, given the rich number of distinct products and the commonality of multi-product firms in the data, such a matching is not practical. Moreover, in small samples, these within-pair comparisons may be noisy, making it difficult to obtain precise causal estimates.

³⁷We have also conducted the primary analyses under alternative weighting assumptions—for example, letting $\hat{\Delta}_p$ be an average where countries are weighted by their total global share of imports of p . The results are similar under alternative specifications.

³⁸If firms in the sample that export better-shocked products also happen to export to the subset of destinations with better shocks for their respective products, then the two components may be correlated.

think of these shocks as being generally negative relative to firms' expectations, with few truly advantageous shocks.

4.3 Estimating Equations

4.3.1 Firm-Level Analysis

Our analysis proceeds in two parts: First, we study reduced form effects on firm-level outcomes; second, we investigate pass-through to worker-level outcomes. To begin, we implement an event-study design to consider the dynamic, reduced-form effects of the demand shock on firms. Graphical analysis of the dynamic effects facilitates a direct assessment of the timing of the effects on different outcomes, and, by including pre-period observations, this analysis allows for straightforward inspection of different pre-period trends across different levels of the shock variable. For each outcome Y_{jt} we estimate the following dynamic difference-in-differences equation:

$$Y_{jt} = \alpha_t + \delta_j + \sum_{k \neq 2007} \beta_k S_j \times \mathbf{1}\{t = k\} + \sum_{k \neq 2007} \gamma_k X_j^{pre} \times \mathbf{1}\{t = k\} + v_{jt} \quad (13)$$

The event-study design is formalized in the inclusion of the firm fixed effect δ_j , which absorbs any differences in outcome levels in the year 2007. The year-specific coefficients β_t identifies the effects of a one-unit higher shock on the *difference* in the outcomes between 2007 and year k ; the 2007 effect is excluded, as it is absorbed in the fixed effect. We cluster standard errors at the firm level to account for any potential serial correlation in the outcome variables. We then plot the year-specific β_t to implement the aforementioned graphical analysis.

The model includes year fixed effects α_t and a firm fixed effect δ_j . For all outcome years, the shock S_j takes the same fixed value, but as $\mathbf{1}\{t = k\}$ is only nonzero during year k , β_k is only identified off outcome variation in year k . We also allow for inclusion of a vector of fixed pre-period controls X_j^{pre} , which are allowed to have year-specific effects γ_k . In the baseline specification, we only control for pre-period export levels to ensure results are not driven by the baseline correlation between the demand shock and the export intensity of the firm, but we consider robustness to additional pre-period controls as well.

To summarize these effects into a single point estimate, we then estimate standard differences-in-differences regressions comparing differential changes of outcome levels from the “pre-period” (2006-2007) to the years after the incidence of the trade decline (2009-2011).

$$Y_{jt} = \alpha_t + \delta_j + \beta S_j \times Post_t + \sum_{k \neq 2007} \gamma_k X_j^{pre} \times \mathbf{1}\{t = k\} + v_{jt}, \quad t \in \{Pre, Post\} \quad (14)$$

In the benchmark specification $Pre \equiv \{2006, 2007\}$ and $Post \equiv \{2009, 2010, 2011\}$, so as to capture effects that appear either during the 2009-2010 episode or with a one-year delay. While the treatment effect is pooled across years, we continue to control for year-specific effects of pre-period covariates and yearly fixed effects. In this specification, average pre-period

differences (in both pre-period years) across differently shocked firms are absorbed by the firm fixed effect δ_j , and β identifies a pooled effect of a one-unit increase of the shock on the difference in the outcome between the pre-period and post-period.

4.3.2 Worker-Level Analysis

After characterizing firm-level effects, we then turn to analysis of worker-level effects on the incumbent cohort of employees at affected firms in 2007.³⁹ Before estimating pass-through elasticities, we consider the reduced-form effects on wages of individuals in this cohort. We implement versions of the dynamic difference-in-differences regression in (13) and the pooled difference-in-differences regression in (14), where the outcome is the average logged wage taken across all individual full-time incumbent workers at affected firm j in 2007. By studying effects on a fixed cohort rather than contemporaneous employees, we rule out the possibility that wage changes are driven by changes in the composition of workers employed by the affected firm. Moreover, we study outcomes for workers regardless of whether they remain at the incumbent firm, which permits identification of changes in rents that are realized by moving out of the treated firm.⁴⁰ Though we implement these regressions at the cohort level, we weight observations by the number of cohort members so that the analysis is approximately the same as a worker-level regression with a cohort-level treatment.⁴¹ In our primary specification, we restrict our focus to the cohort of attached incumbents, which we define as those individuals who were observed working full-time at the affected firm j at each of 2005, 2006, and 2007, as these workers are more likely to remain at the firm in the future in the absence of any intervention.

To quantify the magnitude of these effects, we estimate a pass-through elasticity of wages with respect to sales, which measures the causal change in wages per one-percent change in sales that is due solely to the demand shock. Formally, we estimate the elasticity $\epsilon^{w,R}$ of incumbent wages with respect to observed output R_j by estimating the following difference-in-differences specification, analogous to (14), using the demand shock S_j as an instrument for output R_j , which is the independent variable:

$$\bar{w}_{jt}^{incumb} = \alpha_t + \delta_j + \epsilon^{w,R} R_j \times Post_t + \sum_{k \neq 2007} \gamma_k X_j^{pre} \times \mathbf{1}\{t = k\} + v_{jt}, \quad t \in \{Pre, Post\} \quad (15)$$

where here j denotes the cohort employed at the affected firm in 2007. Formally, we instrument for $R_j * \mathbf{1}\{t \in Post\}$ with $S_j * \mathbf{1}\{t \in Post\}$ using a two-stage least squares estimator.⁴² When

³⁹Our worker-level analyses only include worker years where individuals are employed full time (≥ 30 hours per week), both in determination of the incumbent cohort and for measurement of outcomes.

⁴⁰We consider both specifications that include worker outcomes at any firm in the post period and specifications where only workers remaining at the 2007 are considered.

⁴¹When workers do not appear in the data in post-period years, they are not included in the cohort-average outcome in that observation year, thus the underlying sample of workers is not balanced. However, we consider specifications where missing observations are treated as zeros and incorporated into the average; results are not sensitive to this adjustment. We fix observation weights over years as the number of full time workers in the 2007 cohort.

⁴²Since all pre-period variation correlated with the treatment is absorbed in the firm fixed effect, it is only necessary

the exclusion restriction holds, $\epsilon^{w,R}$ identifies the elasticity by which exogenous, idiosyncratic changes to firm output or productivity pass through to incumbent wages. We use both sales and value added as alternative measures of output R_j ; in practice, the results are not sensitive to the choice of output measure. For comparison, we also estimate (15) by “naive” OLS on the same data, without using S_j as a shock.

We also test for differential effects among specific subgroups of employees by implementing (15) above, taking the average employee outcome only among workers in the specified subgroup. Whenever we study effects on that subgroup, we weight cohort observations by the number of 2007 incumbent employees in the specified subgroup.

4.3.3 Heterogenous Pass-Through

In order to test whether pass-through effects are present in all types of firms or only in select subsamples of theoretical interest, we also estimate variants of (13), (14), and (15) that allow treatment effects to differ for high relationship-durability and low relationship-durability industries. Letting H_j denote an indicator for above-median (“high”) industrial relationship durability and $L_j = 1 - H_j$ denote the corresponding “low” indicator, we interact the treatment with *both* H_j and L_j and omit the main effect of the treatment, always controlling for the main effect of H_j in interacted specifications. The interacted event-study specification is

$$\bar{w}_{jt}^{incumb} = \alpha_t + \delta_j + \sum_{k \neq 2007} \beta_k^H S_j \times H_j \times \mathbf{1}\{t = k\} + \sum_{k \neq 2007} \beta_k^L S_j \times L_j \times \mathbf{1}\{t = k\} + \sum_{k \neq 2007} \gamma_k^H H_j \times \mathbf{1}\{t = k\} + \sum_{k \neq 2007} \gamma_k X_j^{pre} \times \mathbf{1}\{t = k\} + v_{jt} \quad (16)$$

The pooled reduced-form difference-in-differences equation is analogous.

When estimating heterogeneity pass-through elasticities, it is important to account for the possibility that the first stage effects of the shock differs across subsamples. Thus it is necessary to separately instrument for R_j in each subsample. When estimating heterogeneity in the pass-through elasticity, we estimate the interacted equation

$$\bar{w}_{jt}^{incumb} = \alpha_t + \delta_j + \epsilon_H^{w,R} R_j \times Post_t \times H_j + \epsilon_L^{w,R} R_j \times Post_t \times L_j + \gamma^H \times Post_t \times H_j + \sum_{k \neq 2007} \gamma_k X_j^{pre} \times \mathbf{1}\{t = k\} + v_{jt} \quad t \in \{Pre, Post\} \quad (17)$$

With two endogenous regressors, it is necessary to separately instrument for the shock effect among high- and low-durability firms. Thus, we instrument for $R_j \times Post_t \times H_j$ using $S_j \times Post_t \times H_j$ and likewise for the interaction with L_j .

4.4 Identification Assumptions

The key identification assumption is that differential changes in the outcome are only due to differential levels of the shock variable, which requires that all unobserved determinants of the

to instrument for the interaction with the post period

outcome (reflected in v_{jt}) evolve in parallel for any level of the shock S_j , conditional on observables. This is the standard difference-in-differences “common trends” assumption. Formally, the shock must be uncorrelated with all unobserved determinants of changes in firm performance or wages since the pre-period: for the dynamic specifications, it must be that $E[S_j \times \Delta v_{jt}] = 0 \forall t$, where Δ denotes the difference in year t relative to 2007; for the pooled specification, the equivalent condition is that $E[S_j \times Post_t \times v_{jt} | \delta_j] = 0 \forall t$. To identify valid pass-through elasticities based on an endogenous output measure R_j it must be the case that all variation in R_j correlated with the shock S_j is due only to the same underlying idiosyncratic, exogenous variation in demand that causes the change in wages, conditional on fixed effects and controls. Put differently, the orthogonality assumptions must jointly hold both for output and wage outcomes.

The conceptual framework highlighted two types of violations are of particular concern: endogeneity of shocks to labor supply factors (failures of the exogeneity assumption) and correlation of shocks with market-level demand factors (failures of the idiosyncrasy assumption). The estimates will be biased if the shock is not exogenous to worker inputs—in particular, if firms with better latent trends in firm performance and worker wages systematically sort to destinations with better demand shocks. The conjecture that estimates simply reflect differential trends in latent drivers of wage growth or output growth, which would occur regardless of the demand shock, is partially testable. If results are truly driven by the shock, rather than unobserved trend heterogeneity, then one should not observe differential evolutions of outcomes associated with the demand shock *prior* to the onset of the recession. We provide graphical evidence that the demand shocks only predict changes in firm behavior after the onset of the recession by plotting the year-specific effects of the demand shock on the primary outcomes obtained from the event-study specification in (13).

Even in the absence of differential pre-period trends, exogeneity could still be violated if firms and workers with a higher propensity to fare better during the recession (higher resilience) systematically sorted to destinations markets that were less affected by the recession. One concern is that a large share of Portuguese exports go to neighboring Spain, which experienced a particularly adverse recession in the same years. Since the decision to export to Spain is surely different than decisions to export to other destinations given its immediate proximity to Portugal, worse shocks due to exposure to Spain may be reflective of latent characteristics of firms and workers. Similarly, Angola, Portugal’s largest former colony, is a top trading partner that experienced particular strong growth during 2009—thus, good shocks reflecting exposure to Angola may also be reflective of different latent characteristics of firms and workers. Thus, in all the primary analyses, we control for year-specific effects of the total pre-period share of exports that go to Spain and to Angola respectively.⁴³ We also consider robustness of results to controls for year-specific effects of exposure of each of the other of Portugal’s top ten destination countries.

⁴³The results are not sensitive to the inclusion or exclusion of these controls, in practice.

Several additional falsification tests provide additional evidence that firms and workers do not sort across destinations based on latent “resilience” (i.e., a tendency to perform better specifically after the onset of the Great Recession). First, one can test whether it is plausible that better-improving firms *strategically* sorted across destination markets by testing whether changes in demand during the recession were highly forecastable. Appendix Figure A.2 shows that there is almost zero correlation between change in imports in country-by-product markets during the recession and during the change that occurred in the three years prior; a similar finding occurs if one solely considers changes in demand across country, conditional on product fixed effects.

Second, one can use the detailed relationship-level data on firms with multiple destinations to directly test whether firms that had better across-the-board export performance selected into relationships with customers in foreign markets that had better demand shocks. Following Khwaja and Mian (2008), we estimate relationship-level regressions of firm exports to a foreign markets on the import demand shock in the destination. If firms with multiple export destinations do not sort on latent firm-level drivers of export performance, the estimates from this regression should be invariant to controls for firm fixed effects, which absorb firm-level drivers of export performance. In Appendix B.1, we discuss this test in greater detail and present results consistent with no sorting on unobservables.

Third, although differences in baseline levels of the outcomes and covariates across differently shocked firms do not in themselves violate the identification assumptions, it would be suspect if demand shocks were systematically correlated with many important firm productivity and worker characteristics. Figure 3 shows the baseline correlation of the idiosyncratic demand shock S_j with 2007 attributes of the firm. The shock is uncorrelated with all measures of firm size, firm productivity, and employee characteristics; however, there is a significant correlation with the firm’s pre-period export level. Exogeneity would be violated if different export levels are correlated with latent drivers of firm performance and wage outcomes, even while total sales, firm productivity, and employee characteristics are not; to ensure such a correlation does not drive the results, we control for year-specific effects of exports (in logs, levels, and as a share of sales) in the primary analysis. Additionally, we show that results are robust to controls for a wide range of pre-period covariates and industry-level fixed effects.

The other primary threat to identification of true wage pass-through effects is the possibility that our demand shocks are still correlated with market-level shocks, which shift incumbent workers’ outside options. In this case, the estimated coefficient would reflect incidence of a *market-level* demand shock, without necessarily indicating that firms-specific shocks have incidence on wages. We directly test for market-level effects by testing whether the shock to firm j has effects on the performance or labor demand of other similar firms, which are most likely to both factor into the outside option of j ’s employees and be affected by firm j ’s product demand shock. We find that the shocks do not predict outcomes at other similar firms using multiple alternative definitions of labor markets based on industry and geography; we discuss these tests

in more detail in section 5.3. Another way in which our export shocks could be correlated with market-level conditions is if they predicted domestic demand for sales of the firm's products, which would likely affect all other Portuguese producers of the product. We show that the effects of shocks on sales are mostly accounted for by changes in exports, as opposed to domestic sales.⁴⁴

5 Effects of Demand Shocks on Firm Performance and Labor Demand

This section presents the effects of demand shocks on firms' revenues, output, and labor demand. First, we discuss the demand shifter's effects on actual exports, sales, and production at firms and establish an effect on sales that is largely accounted for by changes in exports. Next, we show that there is a reduced-form effect on total labor demand at shocked firms. We find that employment changes only occur on the hiring margin; there is no baseline effect on retention of incumbents. Finally, we test for a relationship between a firm's shock and the output or labor demand of other firms in the same narrow industry and region. We find no such relationship exists, supporting the validity of our idiosyncratic demand shock.

5.1 Effects on Output

Table 3 presents the baseline effects of the idiosyncratic shock on firms' revenue and production. Firms do not systematically exit in response to the shock; nonetheless 1,278 firms exit the data in at least one of the post-period years. Table 3 presents results both for the full sample, treated as an imbalanced panel except in specifications studying growth rates that treat missing observations as zeros, and for the balanced panel of firms that never exit the data. Appendix Table A.1 shows that the composition of firms and workers in both samples are highly similar. While effects are very similar in both the balanced and imbalanced panels, the results are more precise in the balanced sample.

The export demand shock S_j is a powerful predictor of firm export performance. A predicted ten percent change in exports is associated with a five percent actual change from the pre-period to the post period.⁴⁵ This change is directly reflected in the total sales of the firm, which change by about two percent in response to a ten percent shock. While our shock is a shifter to export demand, rather than total demand for firm sales, the magnitude of the effect on total sales is approximately what one would expect given the sample average pre-period share of sales in exports of 34 percent. One could redefine our independent variable as a direct shock to total sales demand by interacting S_j with each firm's pre-period share of sales in exports. Although we do not adopt this approach in our baseline analysis because firms' export intensities are not randomly assigned, it is nonetheless useful to verify that a shock defined as a one percent change

⁴⁴Berman et al. (2015) argue that some spillover to domestic sales should be expected if shocks affect firms' ability to invest in future growth. Since the export data comes from administrative sources that are subject to reporting thresholds, exports effects may be understated relative to sales effects.

⁴⁵Many firms remain in the data but cease exporting, thus the "intensive-margin" coefficient of .5 is attenuated by potential misspecification of zeros which are treated as missing in the log specifications.

in sales actually shifts sales by one percent. Reassuringly, when the shock is directly interacted with the share of pre-period sales in exports the coefficient is approximately one. There is a small but statistically insignificant effect on domestic sales.⁴⁶ We find that value added changes in approximately equal proportion to sales.

Figure 4 plots dynamic effects (β_k) of the shock on exports, sales, and value added for the balanced panel of firms always present in the data. To facilitate direct comparisons of the magnitudes of these effects, we measure in common units (Euros relative to the pre-period asset stock). For all studied quantities, there is no differential trend across shock levels until the onset of the Recession in 2008. After 2008, recorded exports adjust in response to the shock, and sales adjust roughly Euro-for-Euro.⁴⁷ The proportional change in value added constitutes a smaller change in terms of Euros, as intermediates account for a substantial proportion of the value of sales.

5.2 Effects on Labor Demand and Employment Adjustment

Next, we study how changes in demand for firms' output affects firms' demand for labor—both for incumbent workers and for external hires. Table 4 presents estimates of the effects on the total payroll and the quantity of labor employed by firms in response to shocks. We find that payroll has a clear response to the change in sales, though the effect on payroll is only about two-thirds the magnitude of the effect on total value added. In response to the change in effective revenue productivity that occurs due to the demand shock, firms adjust the quantity of labor employed, though the change in log employment is smaller than the change in log payroll. There is no major change in the average hours worked per employee, inclusive of overtime, however.

To shed more light on the margins of labor adjustment, we exploit the longitudinal employer-employee data to decompose the percent change in full-time employment (change relative to pre-period full-time employment, inclusive of zeros) into the part of employment growth (or decline) due to increased retention (or separations) of incumbent employees and the part due to increases (or decreases) in the number of additional hires made and retained that contribute to the current employment stock. Formally, we note that changes in employment from 2007 until year k can be mechanically decomposed into the number of hires made since 2007 (who have not already left the firm)⁴⁸ and the number of retentions of 2007 incumbents by year t (reflecting separation effects

⁴⁶Berman et al () argue that due to internal economies of scale, domestic sales should be impacted to some degree by export shocks.

⁴⁷Larger effects occur in later years, though the discrepancy with exports may arise in part due to per-shipment reporting thresholds in the export data that attenuate measured effects as the impacts set in.

⁴⁸Since new employees typically depart firm with positive probability in their first few years, an increase in accumulated hires will mechanically generate an increase in separations due to any fixed attrition rate. To distinguish between separations mechanically tied to hires, we account for departed hires in the calculation of net accumulated retained hires $Hires_{jt}$

including layoffs)”

$$\Delta Emp_{j,t,2007} = \underbrace{\sum_{\tau=2007}^t \text{Accumulated Retained Hires}_{j,\tau}}_{Hires_{j,t}} - \underbrace{\sum_{\tau=2007}^t \text{Incumbent Retentions}_{j,\tau}}_{Seps_{jt}}$$

Using this identity, we can directly decompose the reduced form effect of the shock on the percent change in employees into a hiring effect and a layoffs/retention effect: $\frac{\Delta Emp_{j,t}}{Emp_j^0} = \frac{Hire_{t,j}}{Emp_j^0} + \frac{Seps_{j,t}}{Emp_j^0}$.

Figure 5 plots the dynamic effects of the shock on employment. As with sales, there are no different pre-period trends across firms corresponding to different shock levels; consistent with exogenous assignment, firms adjust employment only after the onset of the recession. Due to the nature of the decomposition, the employment effect in each year is equal to the sum of the separation and hiring effects. We find that the observed adjustment in employment is wholly attributable to differences in hiring behavior—there is a tight zero effect on the departure rate of incumbent employees. Even though most shocks are adverse during this period, there is no evidence that the employment adjustment among continuing firms is due to differential incidence of layoffs; this is consistent with anecdotal evidence that firing costs are prohibitively large in Portugal. It is not clear, however, whether the absence of any effect on incumbent labor quantities is in fact due to firing constraints; if such constraints exist, they may not be binding, so long as the baseline departure rate of incumbents over the study period of 17 is sufficiently high to enable firms to implement a desired downward adjustment. More generally, the finding that all employment effects occur on the hiring margin is consistent with models where incumbent labor can not be perfectly substituted with external hires.

5.3 Effects on Output and Labor Demand at Other Firms

The effects of the demand shifter S_j are consistent with identification of an exogenous to shock to demand. Yet, prior to studying wage pass through it is important to first assess whether it appears to be *idiosyncratic*. In particular, identification of pass-through of firm shocks separately from labor market shocks requires that the shocks to the firm are not correlated with changes in performance at other firms in the labor market $\bar{\theta}$ that could in turn change market-level labor demand reflected in $\bar{w}$.

If the effects of shock S_j are firm-specific—that is, idiosyncratic relative to the labor market—one should not observe changes in performance or labor demand at the firms that are most likely both to be affected by the shock and to hire employees with the same skill set as those at the shocked firm. Accordingly, we test for effects on sets of *other* firms in the Portuguese labor market, excluding the treated firm, that are likely to satisfy both these conditions. A natural group to look at are other firms in j 's narrow industry—if the shock is idiosyncratic to j , one should not observe any systematic change in output or payroll at those other firms. For each firm

j , we calculate the leave-one-out mean⁴⁹, sales, value added, payroll, and employment among other firms in the same industry and study whether the shocks predict similar effects on these other firms.

Results are presented in Table 5. While the idiosyncratic shock S_j has a clear effect on output and labor demand at the shocked firm, there is no evidence that the shock affects output or labor demand at other firms in the same industry or the same municipality. While there is some evidence of an effect of on sales for other firms in the same five-digit industry *and* the same municipality, there is no significant prediction for value added and no evidence that those firms' labor demand responds comparably to the the treated firm—that is, there is no effect on payroll or employment. Figure 6 plots the dynamic effects on labor demand (shifts in total payroll) at both shocked firms and other firms in the same industry-municipality pair. Although the idiosyncratic shock is not informative about pre-period for shocked firms or other firms, only shocked firms display a payroll response in the main post-period years (2009–2011). In each of these years, the difference between the effect on shocked firms and other similar firms is statistically significant.

By contrast, the common component C_i of the baseline demand predictor Δ_j appears to have effects common to many firms in the market, as apparent in Table 5. Likewise, Figure 6 shows that common shock has similar dynamic effects on the payrolls of both shocked firms and of other similar firms. It should be noted that common effects across similar firms are *not* mechanically implied by the construction of the shock—"commonality" was inferred based solely on foreign import flows, not based on firm export behavior in Portugal. However, the common component of predicted demand strongly predicts increases in sales, output, payroll, and employment at other firms in the same five-digit industry, whether or not these firms were located in the same region.⁵⁰ This finding underscores the importance of isolating the idiosyncratic component of demand variation, as the baseline shock contained variation that was *not* plausibly specific to single firms, but was instead related to changes in payrolls at firms throughout the labor market.

6 Pass-Through Effects to Workers

In this section, we study the incidence of the firm-specific demand shocks on the wages of incumbent workers who worked present full-time at a given firm. First, we present reduced-form evidence on the dynamic effects of the shock on wages. Next, we estimate pass-through elasticities of output changes to wage changes, which measure the percent changes in incumbent worker wages that result from idiosyncratic shocks that change output by a given percentage. We discuss how these can be estimated and identified using an instrumental variables approach, and show that our findings are robust to alternative specifications. The section concludes by studying heterogeneity in pass-through rates across different subgroups of workers.

⁴⁹Outside firms are weighted by pre-period employment to account for their importance in the labor market.

⁵⁰The large magnitude of the effects is due to the common component of demand also strongly predicting *domestic* sales of firms, as global changes in product demand can be reflected in Portuguese demand as well.

6.1 Reduced-Form Effects on Workers Wages

To provide initial evidence on wage incidence, Figure 7 plots the dynamic reduced-form effects of the shocks on the wages of the workers who were incumbent to the firm in 2007. We focus on the subset of attached workers who were employed full-time at the treatment firm in each of 2005, 2006, and 2007; we begin by studying the log monthly base wage of these workers, regardless of where they were employed.⁵¹ As evidenced in the figure, wages respond following demand shocks, with no differential trend apparent prior to the shock. While these effects on the monthly wage could be driven in theory by changes in hours, the effects are in practice similar for monthly and hourly wages, as effects on hours per worker are small (see Table 4).

To test whether this effect is being driven by workers who switch firms and receive differential rents elsewhere, Figure 7 also plots the effect on the monthly contract wage limiting the sample to observations where the worker remains at the incumbent firm. The effects are nearly identical, regardless of this restriction, suggesting that the primary effects occur within-spells and not due to “wage scars” of job switchers. Given the lack of a layoff effect, this is to be expected. In years during which workers are not present in the QP, they are omitted from the analysis rather than treated as zeros—an omission unlikely to bias the results, as there is no observed effect on non-employment and the findings are robust compared to other scalings of the outcome that incorporate zeros.⁵²

These findings suggest that over the two-to-four year horizon, wages are not so inflexible to prohibit any wage adjustment. While downward nominal wage adjustments are rare in Portugal and generally prohibited (Martins, 2014), it does not appear that downward nominal rigidities were binding during this period, as average wage growth was approximately nine percent on average from 2006–2007 to 2009–2011. To test whether particularly negative shocks had muted wage effects due to binding rigidities, Figure A.3 presents a residualized binned scatter plot of $\text{wage} \times \text{post}$ again the shock, where the residuals are obtained from the specification (14) so that the average slope exactly corresponds to the estimated coefficient. As we residualize both the shock and the outcome on observable controls, this exercise tests for asymmetries in shock responses relative to forecastable differences in shock levels. We include a plot of the quadratic fit to provide graphical evidence of any possible nonlinearities in the causal relationship. The relationship is approximately linear. We have also examined spline regressions where slopes are allowed to differ for firms with positive and negative *absolute* S_j values—if anything, we find that wages are more responsive to negative S_j shocks than to positive shocks. These findings suggest that downward rigidities do not bind for negative values of the shock.

⁵¹Since wages are only observed in a single reference month, we exclude overtime and fringe payments that may vary throughout the year. We begin with the monthly salary as this is the object most likely to remain constant across months, though we show robustness to studying hourly wages and total salaries in the following subsection. Workers who exit the data in a given year are treated as missing in that year, though workers whose initial firm exits the data, but who themselves find employment at other firms in subsequent years are included.

⁵²Even in the absence of an average worker-exit effect, it may still be the case that the change in incumbent wage in part reflects changes in the *composition* of separations across the wage distribution within the firm. However, we find no evidence of differential non-employment effects across wage ranks.

The effect of idiosyncratic demand shocks on incumbent wages that we find cannot be reconciled with fully competitive wage determination. In the framework presented in Section 2, this effect implies the presence of a market friction or a worker-replacement cost that generates a quasi-rent within the employment relationship. However, the reduced-form findings alone do not provide a clear sense of the magnitude of the effects. The next section presents results using an instrumental variables estimator to gauge how large these effects are relative to overall change in performance of the firm.

6.2 Estimates of the Pass-Through Elasticity

In this section, we estimate the magnitude of the incidence of firm-specific demand shocks on workers. Ideally, one would observe the instantaneous demand P_j ; in this case, one could estimate the first-stage effect of the demand predictor S_j on P_j and then estimate the causal relationship between P_j and incumbent wages, using S_j as an instrumental variable for P_j . However, in practice, one can only measure observable behaviors of the firm, such as output. Therefore, we estimate the log change in wages that corresponds to a one-percent change in output or revenues that result from a given firm-specific demand shock. Under the assumption that the elasticity of output with respect to P_j is greater than one, this pass-through elasticity is a lower bound of the true elasticity of wages with respect to P_j itself.

The benchmark estimates of the wage elasticity of attached incumbent workers with respect to output are presented in Table 6. We find an elasticity of approximately 0.15—with estimates generally ranging from 0.13 to 0.16 across specifications—which holds regardless of whether the outcome is the monthly base wage, the hourly wage, or total monthly compensation including overtime and fringe payments; it also holds regardless whether we use sales or value added as the measure of output.⁵³ The elasticity is also similar whether or not one considers all observed job years, including at other subsequent jobs, or whether one limits analysis to within-job-spell changes. Since the baseline effects on separations are negligible, alternative measures of wages that include zeros—for example, the logarithm of one plus the wage—yield similar results.

These estimates are an order of magnitude larger than the coefficients estimated by simple OLS incorporating all observed changes in output, rather than those identified by the demand shock. Across all specifications, we find a small but very precise observed elasticity of about 0.02. Despite concerns that OLS estimates would be upward-biased due to simultaneous effects of labor supply shocks on wages and output or due to common shocks, the OLS appears to be significantly downward biased. In Section 8, we relate this finding to a broader literature that has found similar effects, suggesting that the first-order problem with studying observational pass-through elasticities is mismeasurement of demand shocks (or revenue productivity shocks, more generally).

Table 7 presents results from alternative specifications in order to probe the robustness of

⁵³While value added is perhaps the more relevant output measure, the first stage is weaker because value added is a calculated outcome rather than a directly observed behavior like sales.

results. Column 1 displays our benchmark estimates, and estimates with no controls are in Column 2. Column 3 includes controls for year-specific effects of a rich set of pre-period covariates, and Column 4 includes controls for year-specific effects of exposure to each of the top ten destination countries and for pre-period demand growth at each firm’s destination countries. The sign, magnitude, and statistical significance of the estimates are largely robust to the choice of controls. Column 5 controls for five-digit-industry-by-year fixed effects, so that all effects are identified only off of differential shocks within narrow industries. While this reduces the precision of the estimates, the point estimates remained unchanged. Column 6 limits the sample to manufacturing firms; again, while the estimates identified off the smaller sample are less precise, the point estimates remain similar. To relate the results to the firm-level effects in Tables 4 and 5, Column 7 restricts the sample to the balanced panel of firms that never exit. This restriction has very little effect on the point estimates or their statistical significance. The remaining columns explore robustness to inclusion of the largest firms under alternative weightings. When the baseline specification is estimated using firm-level weights, the magnitudes are similar but somewhat less precise.⁵⁴ When firms are employment weighted, no effect is present when large firms are included, as the results are dominated by noisy zero effects for large firms. However, when these firms are included and firms are weighted equally, the results remain similar regardless of whether the largest firms are included. Thus, while the precision of the estimates varies across specifications, the finding of a pass-through elasticity of 0.15 among the primary sample of small and medium firms exposed to idiosyncratic trade risk is not sensitive to specification.

6.3 Differential Pass-Through to Subgroups of Workers

Previous work has suggested rent-sharing may differ across gender groups [Card et al. \(2016a\)](#) or across income subgroups [Kline et al. \(2017\)](#). Accordingly, we estimate separate pass-through estimates on a number of subgroups of incumbent employees. In each case, we estimate equation (15) where the outcome is the average cohort wage for members in the subgroup, and regressions are weighted by the number of members of the subgroup at the firm in the incumbent cohort. As all subgroups are not equally represented at all firms, any differences across groups may be due to differences in firm characteristics, and, if workers sort across heterogeneous firms, then any worker-type effect heterogeneity may be incidental to the worker’s type. Nonetheless, it is of interest whether any such heterogeneity is apparent.

In practice, we do not find significant heterogeneity in the magnitude of effects across subgroups of workers. Results are presented in Table 8. The pass-through affects onto less-attached workers with less than three years’ tenure and onto workers with fixed-term contracts—regardless of tenure length—are substantially less precise than the effects estimated for workers with longer tenures or permanent contracts. However, the magnitudes of

⁵⁴While the first-stage power actually improves under this alternate weighting, the reduced form effects on wages are noisier when smaller firms are given higher weight.

pass-through for these less-attached workers, with their higher separation hazards, are not significantly smaller. Among attached workers, there are no clear differences in effects across tenure lengths, gender, or pay levels.

7 Heterogeneity in Pass-Through: The Role of Relationship-Specific Surplus

The previous section established that idiosyncratic firm shocks have significant incidence on the wages of workers. In the conceptual framework of Section 2, such pass-through effects occur because firms cannot perfectly substitute for incumbent workers with external hires. If this account is accurate, then pass-through effects should only exist to the extents that labor markets are not fully fluid and that replacing workers is costly. This section studies whether pass-through effects differ across firms that are in more or less fluid labor markets.

Although the true replacement costs of incumbents within firms are not directly observable, theory points to useful proxies. In many models featuring relationship-specific surplus, labor markets with greater replacement frictions are marked by more durable firm-employee relationships (“relationship durability”): once a worker is deemed a good fit, they are less likely to leave the firm than if labor markets were fully fluid. For example, when firms have sunk costly investment in firm-specific human capital that is less useful on the outside market (Becker, 1962; Jovanovic, 1979a; Lazear, 2009), workers and firms have an *ex post* incentive to maintain the employment relationship, as both firm and worker lose the benefit of that investment in the outside labor market. Similarly, when *ex ante* unobservable match quality varies substantially across firms and potential employees such that both parties undergo costly searches to find a good fit (Jovanovic, 1979b; Mortensen and Pissarides, 1994), both parties also have substantial option value to maintaining a relationship once a good match is made. And when firing is costly for institutional reasons, these constraints on downward adjustment are more likely to bind when the baseline attrition rate is lower. Accordingly, we study heterogeneous pass-through across observable measures of relationship durability.

We measure two proxies of relationship durability: typical tenure lengths and separation rates in the detailed industry level. In particular, we focus on separations and tenures of workers who have been offered permanent contracts, as the offer of a permanent contract indicates the firms’ belief that the match quality is high, and as the high barriers to layoffs suggests most separations are due to voluntary quits for these workers. To avoid incorporating mean-reverting behaviors of sample firms into the index, we calculate the typical tenure (weighted average across firms of median permanent employee tenure in 2003–2007) and the separation rate (average ratio of separations to stock of permanent workers across in 2003–2007) for permanent contract workers among all non-primary-sample firms in the five-digit sectors.⁵⁵ We then divide

⁵⁵To calculate the degrees of fluidity most likely to characterize the sample firms, we calculate these averages for all firms with 100 employees or fewer within the industry. These industry-level indicators are highly predictive of the

the analysis sample into two equally sized subsamples for each separate index that correspond to above-sample-median and below-sample-median levels of a specified industry relationship-durability index.

To provide graphical evidence of the differential dynamic wage adjustments across the sub-sectors, Figure 8 plots the coefficients reflecting the year-specific effects of the shock for each subgroup, using both proxies for relationship-durability. The estimates are obtained from an interacted version of the specification plotted in Figure 7—for each year, we estimate and plot year-specific coefficients on the shock S_j , interacted with either an indicator for high durability or for low durability (no main effect is included). In all interacted specifications, we control for year-specific effects of the durability indicator itself. Regardless of the choice of durability index, similar results are found: There is no differential relationship between the shock and wage across durability levels prior to the Recession. However, after 2008, there are *only* significant effects in the sectors with high relationship durability.

While the graphical evidence is informative, a proper test of heterogeneity in *pass-through* requires one to adjust for the possibility that differential wage adjustments may in part reflect differential adjustments in revenues or output. Table 9 presents estimates of the subsample-specific pass-through elasticity using interacted IV specification in equation (17). To identify separate elasticities for high-durability and low-durability sectors, we estimate the interacted differences-in-differences specification using interactions of the shock S_j with each subgroup as instruments for the interactions of output levels with each subgroup, again controlling for the year-specific effects of the subgroup indicators themselves. Figure 9 plots the main estimates for easier comparison.

The results are nonetheless similar, even after accounting for potential difference in the first stage. Pass through in low-durability sectors is never statistically different from zero, and point estimates are generally close to zero. By contrast, the large, precise pass-through effects found in high-durability sectors appear to have been driving the baseline findings in Table 6—and, in some specifications, the point estimates are substantially larger than the pooled estimates, exceeding 0.3. While this contrast is present across specifications, only in the first specification in Table 9 are the estimates precise enough to rule out equality of the elasticities at a five-percent level.

To the extent that these differences in relationship durability reflect differences in frictions and fluidity across labor markets, these finding support our framework, in which positive wage incidence of firm-specific shocks arises due to barriers to replacing workers and jobs. Because these differences across industries are observational, rather than due to some exogenous source of variation in fluidity across sectors, one must be cautious when interpreting the treatment effect of heterogeneity as causal. Nonetheless, there is reason to believe these differences in relationship durability have a direct relationship to frictions that give rise to pay differentials. Appendix Figure A.4 relates our measures of relationship durability to firm wage premiums, estimated as

corresponding tenure lengths and separation rates at the sample firms in the same industry.

AKM firm fixed effects.⁵⁶ We find our measures have a strong cross-sectional correlation with firm pay premiums both at the industry level and at the firm level within industries. Put differently, firms in sectors that appear to have more job or worker replacement frictions pay wages that are higher than can be explained by workers' attributes alone—this supports the claim that these industries have replacement costs that generate larger larger surpluses in the employment relationship. Taken together, these pieces of evidence suggest that firm pay differentials do in fact arise due to pass-through of firm shocks to incumbent employees, particularly in labor markets with higher degrees of imperfections.

8 Discussion

8.1 Interpretation of Effect Magnitudes

We found that idiosyncratic demand shocks that caused a 10 percent change in sales resulted in a roughly 1.5 percent change in the wages of attached incumbent workers—specifically, our main estimates of the pass-through elasticity in Tables 67 ranged between 0.13 and 0.16. Although we estimated this elasticity in a single IV regression, this magnitude roughly corresponds to the reduced-form effect of the demand shock on wages—between 0.03 and 0.04, as in Figure 7—divided by the reduced-form effect on output—approximately 0.20, as in Table 3. In sectors with high relationship durability, measured by the out-of-sample separation rate of workers with permanent contracts, this elasticity was significantly larger and exceeded 0.30 in some specifications.

To assess whether these estimates are economically large, we benchmark them against baseline levels of wage variability in the data. In the analysis sample of incumbent workers, a worker at the 25th percentile experienced a log nominal hourly wage change of 0.033 between 2007 and 2010, while a worker at the 75th percentile experienced a change of 0.163 (the IQR is thus 0.13).⁵⁷ In comparison, the variation in the idiosyncratic demand shock accounts for very little of cross-individual differences in wage growth. For example, the reduced-form estimate of 0.03 implies that moving an individual from a firm with a 25th percentile idiosyncratic export demand shock S_j to a firm with a 75th percentile shock would only increase wages by one percent. Of course, as the variation idiosyncratic shock can only account for a very small portion of overall dispersion in output growth, this small figure is a more direct reflection of the size of the shock than of the size of the effect.

A more useful approach is to consider how much of the observed variation in wage growth can be explained by the *total* variation in output growth, given our pass-through elasticity. This exercise more directly addresses how much of all changes in wages could plausibly be due to firm-

⁵⁶We estimate wage equations of the form $w_{ijt} = \alpha_i + \phi_j + \beta X_{it} + \delta_t + \epsilon_{ijt}$ as in Card et al. (2016a) and Card et al. (2016b), where the firm fixed effects of interest (ϕ_j) are estimated off of workers who move jobs.

⁵⁷The median was 0.08 and the mean was .10. Due to extreme values in the tails, the standard deviation in the raw data (0.179) significantly larger than one half of the IQR. The distribution of wage changes is very similar in both the high-durability and low-durability subsamples.

level factors. The answer is striking: the pass-through elasticity for value added (0.158) implies that moving an individual from a firm at the 25th percentile of log value added growth to a firm at the 75th percentile would increase wages by 6.8 percent—over *half* of the total wage-growth IQR. In fact, limiting the exercise to high-durability sectors, moving an individual from the 25th to the 75th percentile of firm value added growth would increase wages by amount that equal to the full wage growth IQR.

This latter figure implies that, for most workers in the high-attachment sample, dispersion in firm performance on its own could generate as much dispersion in wage growth as is actually observed in practice. Of course, this calculation rests on extreme assumptions that are not realistic: It presumes that *all* variation in value added growth is due to exogenous shocks to firm revenue productivity. Moreover, it assumes that all shocks pass through with the same elasticity we estimated above. Thus, while we find this exercise useful to interpret the magnitude of our results, we are hesitant to make claims about overall wage dispersion based on extrapolations of our results.

8.2 Comparison of Elasticities Using Different Sources of Output Variation

We found that instrumented estimates of the pass-through elasticity are significantly larger than coefficients from OLS regressions of wages on output. A standard concern with observational estimates is the potential that positive relationships between output growth and wage growth reflect changes in worker inputs or changes in market-level demand that shift workers outside options, as opposed to idiosyncratic firm shocks. In that case, the observed coefficients would overstate the causal pass-through of firm-level shocks to workers. Given this concern, our findings are somewhat surprising.

We argue this is because observed changes in output are poor indicators of underlying product market conditions that determine labor demand, leading to attenuation of estimates of demand shocks pass-through. This story is supported by prior observational studies that find the relationship between output and wages nearly doubles in size when changes in sales over a given horizon are instrumented for with changes in sales over a longer horizon to account for pure measurement error. Such findings are consistent with attenuation bias in the OLS, suggests that the independent variable measures the underlying labor demand conditions with substantial error (Card et al., 2016b). Moreover, the problem attenuation due to mismeasurement of revenue productivity shocks is not unique to studies of firm-level pass-through; this problem would similarly attenuate estimates of market-level wage incidence. Indeed, Abowd and Lemieux (1993) found evidence of similar attenuation in their study of industry-level rent-sharing. Similar to our results for firms, they found a trivially small observational relationship between wages and industry value added, but they saw significantly larger wage incidence when industry labor demand was instrumented using the product demand shocks⁵⁸. Although the setting is different, the reason for attenuation may be similar; specifically, observational changes in sales are likely

⁵⁸Card et al. (2016b) show that their estimates imply a productivity pass-through elasticity of 22 percent.

very poor measures of the surplus value in the relationship. This is not to say that changes in sales do not reflect substantive changes in firm conditions; rather, firms may experience many kinds of shocks that do not directly shock demand for incumbent labor. For example, planned firm expansions, increased performance of other factors of production, and foreseen variation in revenues for long-run projects could all create large variations in firm performance that do not discretely increase the demand for incumbent labor at the same time.

We note that, even given significant attenuation bias, OLS estimates may be still further biased due to other confounds. In fact, it may be the case that observational changes in sales similarly attenuate these biases, as observed changes in output measure all components of incumbent revenue productivity with substantial error. Even if OLS estimates are attenuated due to measurement error, pass-through estimates instrumented with market-level demand shocks can still cause upwards bias in estimates of rent-sharing within the firm. Market-level shocks not only affect the internal labor demand for workers; they may also cause wages to rise due to increases of the outside option. One would therefore expect the pass-through elasticities estimated in the common product-level demand to be larger than those estimated off of idiosyncratic shocks.

In practice, we find that firm responses to common shocks C_j and idiosyncratic shocks S_j during the Great Recession were qualitatively different, rather than simply quantitatively different. Appendix Table A.3 shows that for firms that do not exit, the pass-through effects of sales to wages estimated from common shocks are similar to the baseline effects in Table 6. However, in contrast to our findings that idiosyncratic shocks have zero effect on firm exit and incumbent job loss, Appendix Figure A.5 shows that common shocks have large effects both on firm exit and on the probability that 2007 incumbent employees exit the data completely.

It is therefore difficult to directly compare the magnitude of the impacts of common shocks to the idiosyncratic shocks, as one cannot observe the underlying demand (or “potential sales”) level for adversely shocked firms that exit the data. Because the common shock, unlike the idiosyncratic, leads firms to exit—perhaps selectively—results vary dramatically depending on how one conditions on firm survival. Moreover, Appendix Figure A.5 shows that, in contrast to our findings based on idiosyncratic shocks, continuing firms are more likely to adjust *quantities* of incumbents in response to product-level shocks in the long run. When firms adjust quantities of incumbents by making layoffs, wage effects for remaining incumbents may be muted, as per-remaining-incumbent surplus changes less. Further, if firms selectively lay off workers based on unobserved attributes, observed wage effects on remaining incumbents may be further muted by selection bias. Finally, it is important to note that product-level demand shocks result in persistently larger changes in sales—in contrast the findings of Section 6, downward nominal wage rigidities more plausibly begin to bind, further muting the response. Thus, while layoffs and nominal rigidities were not first-order issues when studying the idiosyncratic shocks in our particular settings, they may be important in other situations and should be accounted for when estimating rent-sharing effects.

8.3 Comparison to Previous Literature

Similar to our OLS results, prior studies have typically found small but very robust relationships between observed changes in sales or production and observed changes in wages or earnings in longitudinal employer-employee matched datasets; Manning (2011) and Card et al. (2016b) provide detailed surveys of this work. In Italian data similar to ours, Card et al. (2013a) find a .04 longitudinal elasticity of wages to output after adjusting for changes in outside options of workers. Guiso et al. (2005) argue that the small magnitude of observed relationship between short-run changes in sales and wages may be partially explained if risk-neutral firms insure risk-averse workers against transitory shocks, as in Baily (1974), Azariadis (1975) and Akerlof and Hajime Miyazaki (1980). In the same data, they find that wages appear to be invariant to transitory shocks—as inferred based on the fit of a structural time series process—but sensitive to permanent shocks to firm income, though only with an elasticity of 0.06.⁵⁹ Similarly, Card et al. (2016b) find that in the population QP data, longitudinal correlations between firm performance and employees’ wages yield small elasticities of (0.06 or less) even over five year horizons.

However, the larger magnitudes of our baseline pass-through estimates—which are a lower bound for the elasticity of incumbent wages with respect to firm productivity—are highly similar to the associations found in cross-sectional comparisons of wage premiums and revenue productivity. In particular, studying the same employer-employee data in Portugal during the period 2005–2009, Card et al. (2016b) find that the coefficient found when AKM firm fixed effects are regressed on value added per worker across firms is approximately 0.13—much larger than their observed longitudinal correlation between firm output changes and wages but remarkably similar to our baseline estimates of approximately 0.15.⁶⁰ The magnitude of our point estimates is large enough to account for a causal relationship between firm pay differentials firm output levels of the magnitude they observe.

It should be noted, however, that firm pay premiums identified through firm movers and those identified through pass-through of shocks to incumbents may reflect distinct types of quasi-rent-sharing. Estimates of firm pay premiums obtained from studies of job-movers reflect pay changes that are realized quickly upon joining a firm—and that may or may not be invariant to firm productivity. By contrast, our framework emphasizes rent-sharing in internal labor markets, where incumbent wages may change relative to outside options due to shocks but no premiums accrue to workers who join the firm after a shock is realized. Relatedly, we note that even if our results can account for a large disparities in wage growth over a medium-run horizon, it

⁵⁹ A similar result was subsequently found in Portugal by Cardoso and Portela (2009)

⁶⁰ We do not use per-worker measures of output to proxy for revenue productivity *changes* in response to well-defined shocks, as the dynamic response of per-worker measures additionally reflect the behavioral response of firms that adjust labor quantities in response to the shock. Importantly, while one can bound the degree of revenue productivity pass-through using changes in total output, there is *no* clear relationship between changes in per-worker output and underlying revenue productivity due to these quantity responses. Indeed, when firms are Cobb-Douglas producers and perfectly flexible price-takers, Y/L remains constant in response to all changes in revenue productivity. Nonetheless, measures are arguably better proxies for productivity differentials in the observational *cross – section* of firms during a given time period, as these measures reflect static differences in outputs given a fixed number of inputs.

is uncertain whether accumulation of such medium-run changes could generate cross-sectional wage dispersion similar to that observed in practice.⁶¹

These results also relate to findings of changes in wages and firm profitability in response to firm innovations (Van Reenen, 1996) or to unexpected approval of patent applications (Kline et al., 2017). The magnitude of the pass-through effects found here are comparable on magnitude to the effects found in those studies. While demand shocks and grants of intellectual property both boost profitability of firms, however, they also differ somewhat in nature. The finding by Kline et al. (2017) that grants of valuable patents affect earnings shortly after their award, with the largest effects for employees listed on the patent, leads the authors to suggest that these pass-through effects reflect incentive pay arrangements. For example, research staff may be rewarded for past work—which had been of unclear value *ex ante*—after the fruits of the labor were revealed to be valuable. Importantly, although the timing of the award may be quasi-random, the work that went into the patent can be directly attributed to individual workers. Even though the non-contractibility of research quality does, in a sense, constitute a market imperfection that generates wage pass-through, it remains unclear whether these effects imply workers would share incidence of firm demand shocks, especially when demand conditions are perceived to be exogenous to effort. Similarly, it is not clear whether their results imply that one should expect a firm-specific subsidy to have incidence on the wages of incumbent workers. The frictions generating wage pass-through in our setting—which we have argued reflect costs to replacing jobs and workers—appear to be different than the contracting frictions that generates rent-sharing in patenting firms.

8.4 External Validity

While we find significant incidence of idiosyncratic firm demand shocks in Portugal during the Great Recession, the important questions of how behavior would differ in different institutional settings or in response to different kinds of shocks remain. Our finding that pass-through effects only occur in industries with higher levels of relationship durability raises the possibility that rent-sharing behavior may differ substantially in alternative contexts. Our setting, Portugal, is characterized by very strong labor market protections—while many economies feature similar institutional protections, in other advanced economies such as the United States, employment protections are significantly weaker and employment is “at-will.” It therefore may be the case that even the “high-durability” industries would feature less rent-sharing in alternative regulatory contexts.

In addition, our finding that firms engage in qualitatively different behavior in response to large changes in global demand for entire product groups raises the possibility that real-world firms adjust differently to different types of shocks. In particular, firms may respond to differently during expansionary periods to the same idiosyncratic shocks as those studied here, compared to our results taken from the recession. We focus on the Great Recession for purposes

⁶¹The answer depends both on the persistence and mean-reversion of shocks and on the persistence of wage effects.

of identification, as we believe import changes were harder to foresee during this period; we nonetheless believe that it would be useful to understand how unforeseeable demand shocks would affect wages during better times.

9 Conclusion

This paper presented evidence that exogenous, firm-specific shocks have significant incidence on employees' wages. We found that an export demand shock that changes firm output, measured via sales or value added, by 10 percent leads to a 1.5 percent change in the wages of attached incumbent employees. These findings cannot be reconciled with fully competitive wage determination in markets where firms are price-takers. Accordingly, these effects were concentrated in labor markets characterized by more durable employment relationships and lower fluidity, which can indicate larger labor market imperfections, where pass-through effects exceed three percent per 10 percent output change. The pass-through elasticities we estimate with respect to output provide a lower bound for the magnitude of the incidence of the underlying demand shocks on worker's wages.

An important implication of these findings is that *where* one works can make a significant difference in how much a worker is paid, regardless of the skills, abilities, and efforts one brings to a job. Thus, this work adds to a growing body of research that has found that a significant amount of wage dispersion can be attributed to cross-firm pay differentials, even conditional on worker attributes or fixed effects. We provide novel evidence that this observed relationship may indeed be causal in nature in economies with frictional labor markets similar to those in Portugal. A related implication is that, in such labor markets, policies that are designed to enhance the performance of individual firms, such as targeted production or investment subsidies, may have incidence the specific workers employed at targeted firms. The potential for incidence on workers should be taken into account in distributional analyses of supply-side policies that may not be redistributive in intent.

This paper sheds new light on how wages are determined in imperfect markets, which is of direct importance for evaluating the distributional incidence of policies and economic shocks, as well as channels by which shocks propagate throughout economies. In addition to the distributional implications, however, the presence of imperfect competition likely has important consequences for the dynamic efficiency of labor markets since anticipated rent-sharing effects may lead to distortions in hiring and training decisions. While a rich body of theoretical work has characterized how such distortions may occur in principle, an important area for future work will be to build upon empirical studies of wage-pass through and employment adjustment in frictional economies—such as this paper—to empirically quantify how large the inefficiencies are in general equilibrium. In particular, it would be useful for future work to consider what types of social welfare functions could justify introduction of labor market protections that achieve distributional goals at the cost of aggregate efficiency.

References

- Abowd, John A. and Thomas Lemieux**, "The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada," *The Quarterly Journal of Economics*, November 1993, 108 (4), 983–1014.
- , **Francis Kramarz**, and **David Margolis**, "High Wage Workers and High Wage Firms," *Econometrica*, March 1999, 67 (2), 251–333.
- Acemoglu, Daron and William B Hawkins**, "Search with Multi-Worker Firms," *Theoretical Economics*, 2014, 9, 583–628.
- Addison, John, Pedro Portugal, and Hugo Vilares**, "Unions and Collective Bargaining in the Wage of the Great Recession: Evidence from Portugal," *British Journal of Industrial Relations*, 2017, 55 (3), 551–576.
- Akerlof, George and Hajime Miyazaki**, "The Implicit Contract Theory of Unemployment meets the Wage Bill Argument," *The Review of Economic Studies*, January 1980, 47 (3), 321–338.
- Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser**, "Firms and the Decline in Earnings Inequality in Brazil," *American Economic Journal: Macroeconomics*, 2018, 10 (1), 1–43.
- Autor, David H., David Dorn, Gordon Hanson, and Jae Song**, "Trade Adjustment: Worker-Level Evidence," *Quarterly Journal of Economics*, 2014, 129 (4), 1799–1860.
- Azariadis, Costas**, "Implicit Contracts and Underemployment Equilibria," *Journal of Political Economy*, December 1975, 83 (6), 1183–1202.
- Baily, Martin Neil**, "Wages and Employment under Uncertain Demand," *The Review of Economic Studies*, January 1974, 41 (1), 37–50.
- Bank, The World**, "World Bank National Accounts Data," August 2016.
- Barth, Erling, Alex Bryson, James C. Davis, and Richard Freeman**, "It's Where You Work: Increases in Earnings Dispersion across Establishments and Individuals in the U.S.," *Journal of Labor Economics*, 2016, 34 (S2), S67–S97.
- Becker, Gary**, "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, October 1962, 70 (5), 9–49.
- Bentolila, Samuel and Giuseppe Bertola**, "Firing Costs and Labour Demand: How Bad is Eurosclerosis?," *The Review of Economic Studies*, July 1990, 57 (3), 381–402.
- Berman, Nicolas, Antoine Berthou, and Jerome Hericourt**, "Export dynamics and sales at home," *Journal of International Economics*, 2015, 96, 298–310.

- Bertola, Giuseppe and Ricardo Cabellero**, “Cross-Sectional Efficiency and Labour Hoarding in a Matching Model of Unemployment,” *Review of Economic Studies*, 1994, 61, 435–456.
- Bertrand, Marianne and Sendhil Mullainathan**, “Is There Discretion In Wage Setting? A Test Using Takeover Legislation,” *RAND Journal of Economics*, 1999, 30 (3), 545–554.
- and —, “Are CEOs Rewarded for Luck? The Ones Without Principals Are,” *The Quarterly Journal of Economics*, 2001, 116 (3), 901–932.
- Blanchard, Olivier and Lawrence H Summers**, “Hysteresis and the European Unemployment Problem,” *NBER Macroeconomics Annual 1986*, 1986, 1, 15 – 90.
- and **Pedro Portugal**, “What hides behind an unemployment rate: comparing Portuguese and US labor markets,” *American Economic Review*, March 2001, pp. 187–207.
- Braguinsky, Serguey, Lee Branstetter, and Andrew Regateiro**, “The Incredible Shrinking Portuguese Firm,” 2011.
- Brown, James and Orley Ashenfelter**, “Testing the Efficiency of Employment Contracts,” *Journal of Political Economy*, June 1986, 94 (3.2), S40–S87.
- Budd, John, Jozef Konings, and Matthew Slaughter**, “Wages and International Rent Sharing in Multinational Firms,” *Review of Economic Studies*, 2005, 87 (1), 73.
- Card, David, Anna Cardoso, and Patrick Kline**, “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *Quarterly Journal of Economics*, May 2016, 131, 633–686.
- , — , **Joerg Heining, and Patrick Kline**, “Firms and labor market inequality: Evidence and some theory,” *National Bureau of Economic Research*, 2016.
- , **Francesco Devicienti, and Agata Maida**, “Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data,” *The Review of Economic Studies*, 2013, 81 (1), 84–111.
- , **Joerg Heining, and Patrick Kline**, “Workplace Heterogeneity and the Rise of West German Wage Inequality,” *Quarterly Journal of Economics*, August 2013, 128, 967–1015.
- Cardoso, Anna and M Portela**, “Micro foundations for wage flexibility: Wage insurance at the firm level,” *Scandinavian Journal of Economics*, 2009, 111 (1), 29–50.
- Davis, Stephen J., John Haltiwanger, and Scott Schuh**, *Job Creation and Destruction*, Cambridge, MA: MIT Press, 1996.
- Doeringer, Peter and Michael Piore**, *Internal Labor Markets and Manpower Analysis*, Lexington, Mass: D.C. Heath and Company, 1971.
- Dunlop, John**, *Wage Determination Under Trade Unions*, New York, NY: A. M. Kelly, 1950.


- Eaton, Jonathan, Samuel Kortum, Brent Neiman, and John Romalis**, "Trade and the Global Recession," *American Economic Review*, 2016, 106 (11), 3401–2438.
- Foster, Lucia, John Haltiwanger, and Chad Syverson**, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 2008, 98 (1), 394–425.
- Furman, Jason and Peter Orszag**, "A Firm-Level Perspective on the Role of Rents in the Rise in Inequality," in "Presentation at "A Just Society" Centennial Event in Honor of Joseph Stiglitz Columbia University" October 2015.
- Gauiler, Guillaume and Soledad Zignago**, "BACI: International Trade Database at the Product-Level. The 1994-2007 Version," October 2010.
- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi**, "Insurance within the Firm," *Journal of Political Economy*, 2005, 113 (5), 1054–1087.
- Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J. Redding**, "Trade and Inequality: From Theory to Estimation," *The Review of Economic Studies*, January 2017, 84 (1), 357–405.
- Holmstrom, Bengt**, "Moral Hazard and Observability," *The Bell Journal of Economics*, 1979, 10 (1), 74–91.
- Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang**, "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, 2014, 104 (6), 1597–1629.
- Jaeger, Simon**, "How Substitutable are Workers? Evidence from Worker Deaths," December 2015.
- Jovanovic, Boyan**, "Firm-Specific Capital and Turnover," *Journal of Political Economy*, 1979, 87 (6), 1246–60.
- , "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 1979, 87 (5), 972–990.
- Katz, Lawrence F**, "Efficiency Wage Theories: A Partial Evaluation," *NBER Macroeconomics Annual* 1986, 1986, 1, 235–290.
- Khwaja, Asim Ijaz and Atif Mian**, "Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market," *American Economic Review*, 2008, 98, 1413–1442.
- Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar**, "Who Profits from Patents? Rent Sharing at Innovative Firms," 2017.
- Lazear, Edward**, "Job Security Provisions and Employment," *The Quarterly Journal of Economics*, 1990, 105 (3), 699–726.

- , “Firm-Specific Human Capital: A Skill-Weights Approach,” *Journal of Political Economy*, October 2009, 117 (5), 914–940.
- Lindbeck, Assar and Dennis Snower**, “Insiders Versus Outsiders,” *Journal of Economic Perspectives*, 2001, 15 (1), 165–188.
- Manning, Alan**, “Imperfect Competition in the Labor Market,” in “Handbook of Labor Economics,” Vol. Volume 4b, Elsevier, 2011, pp. 973–1041.
- Martins, Pedro**, “30,000 Minimum Wages: The Economic Effects of Collective Bargaining Extensions,” October 2014.
- Mayer, Thierry, Marc Melitz, and Gianmarco Ottaviano**, “Product Mix and Firm Productivity Responses to Trade Competition,” December 2016.
- Mortensen, Dale and Christopher Pissarides**, “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, July 1994, 61 (3), 397–415.
- Oi, Walter**, “Labor as a Quasi-Fixed Factor,” *Journal of Political Economy*, December 1962, 70 (6), 538–555.
- Pissarides, Christopher**, *Equilibrium Unemployment Theory*, 2nd edition ed., Cambridge, MA: MIT Press, 2000.
- Reenen, John Van**, “The Creation and Capture of Rents: Wages and Innovation in a Panel of U. K. Companies,” *The Quarterly Journal of Economics*, February 1996, 111 (1), 195–226.
- Schoefer, Benjamin**, “The Financial Channel of Wage Rigidity,” 2015.
- Solow, Robert**, “Another Possible Source of Wage Stickiness,” *Journal of Macroeconomics*, 1979, 1 (1), 79–82.
- Song, Jae, Nicholas Bloom, David Price, Fatih Guvenen, and Till von Wachter**, “Firming Up Inequality,” *National Bureau of Economic Research*, 2015.
- Stole, Lars A. and Jeffrey Zwiebel**, “Intra-firm Bargaining under Non-binding Contracts,” *Review of Economic Studies*, 1996, 63 (3), 375–410.
- Yagan, Danny**, “Is the Great Recession Really Over? Longitudinal Evidence of Enduring Employment Impacts,” 2016.

Figures and Tables


Figures

Figure 1: Growth and Exports in Portugal Around the Great Recession




Source: World Bank. Figure plots annual GDP and total exports for Portugal in real Euros, with each variable indexed to its 2000 level. Vertical lines indicate the start of the Great Recession in the US at the end of 2007 and the beginning of the Portuguese sovereign debt crisis in the spring of 2011.

Figure 2: Distribution of Demand Predictors and Idiosyncratic Component




Notes: Figure displays kernel density plots, based on an Epanechnikov kernel, of baseline demand predictor Δ_j and the idiosyncratic demand shock S_j . Also displayed is the “pre-period” predictor for firm j , holding its exposure weights fixed as in the calculation of Δ_j , but using the symmetric growth rate from 2003-2004 to 2006-2007 of imports of product p to country c .

Figure 3: Correlation of Shock S_j with 2007 Covariates




Notes: Figure displays standardized correlation coefficients between the 2007 level of the y-axis variable and the idiosyncratic shock S_j . Coefficients and confidence intervals are obtained from a standardized regression with controls for baseline exposure of firms to each of Angola and Spain. Firms are weighted by pre-period average full-time employees, corresponding to the primary analysis.

Figure 4: Year-Specific Effects on Exports, Sales, Value Added, and Payroll Scaled by Pre-Period Firm Assets



Notes: Figure shows year-specific effects of the idiosyncratic shock S_j on balance sheet items in common units (Euros scaled by the affected-firm pre-period asset stock in 2007). Sample is balanced panel of firms that are employ at least one full-time worker in all years, $N = 2,923$. Figure displays year-specific coefficients from regressions of the specified outcome on the interaction between the idiosyncratic demand shock S_j and an indicator for each year, with all interactions estimated jointly as in equation (13). Estimates for each outcome are from separate regressions. Figure plots confidence intervals based on standard errors clustered at the firm level to account for potential serial correlation of errors. Outcomes are all in Euro units (including zeros), and all are scaled by the average 2005-2007 level of assets at the firm. “Exports” are tabulated from the export transaction database, all other outcomes are from the IES database. Regressions are weighted by the average number of full-time employees in 2005, 2006, 2007. All regressions include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level.


Figure 5: Year-Specific Effects on Employment Adjustment: Hiring Versus Incumbent Retention




Notes: Figure shows year-specific effects of the idiosyncratic shock S_i on employment, and fully decomposes the main effect into two margins of adjustment. Sample is balanced panel of firms that are employ at least one full-time worker in all years, $N = 2,923$. Figure displays year-specific coefficients from regressions of the specified outcome on the interaction between the idiosyncratic demand shock S_i and an indicator for each year, with all interactions estimated jointly as in equation (13). Estimates for each outcome are from separate regressions. Figure plots confidence intervals based on standard errors clustered at the firm level to account for potential serial correlation of errors. Outcomes are all tabulated from the employer-employee matched dataset, units are counts (based on full-time workers, including zeros) scaled by 2005-2007 average full-time employment at the firm. “Retention of 2007 Incumbents” is the percentage of 2007 incumbents present at the firm in the baseline year. “Accumulated hires” is the total number of hires made since 2007, less the number of new hires that have left the firm by the observation years, and is algebraically equal to the effect on employment less the effect on retention (prior to 2007 counts of hires are subtracted rather than added so this identity holds in all years). Regressions are weighted by the average number of full-time employees in 2005, 2006, 2007. All regressions include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level.

Figure 6: Payroll Effects at Affected Firms and Other Similar Firms

Year-Specific Effect of Idiosyncratic Component S_j on Log Payroll




Year-Specific Effect of Common Component C_j on Log Payroll



Notes: Figure shows year-specific effects of the idiosyncratic demand component S_j and the common demand component C_j on both the payroll of directly affected firms and on the average payroll of other firms in the same detailed industry and municipality. During the main post-period years (2009–2011), the idiosyncratic component affects labor demand at directly-affected firms, but not other similar firms—however, the common component has similar effects on both the directly-affected firms and on other similar employers. The sample of affected firms is the balanced panel of firms that are employ at least one full-time worker in all years, $N = 2,923$. For this same sample, we calculate the employment-weighted average log payroll of all other firms in the same 5-digit industry and the same municipality (one of 24 subregions in Portugal). Figure displays year-specific coefficients from regressions of the specified outcome on the interaction between the idiosyncratic demand shock S_j and an indicator for each year, with all interactions estimated jointly as in equation (13). Estimates for each type of shock and each set of firms are from separate regressions. For comparability, we employ the same specification as in (4) when studying both own-firm and other-firm effects. Regressions are weighted by the average number of affected-firm full-time employees in 2005, 2006, 2007. All regressions include year fixed effects, as well as controls for year-specific effects of 2005–2007 affected-firm exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level for own-firm effects, and the 5-digit-industry-by-municipality level for the other-firm effects.


Figure 7: Dynamic Effects on Log Contract Wage of Attached Incumbents



Notes: Figure shows year-specific reduced-form effects of the idiosyncratic shock S_j on the wages of individuals who worked full-time at shocked firms in 2005-2007. Sample is all firms in analysis sample with at least one attached incumbent, defined as workers who were present at the treated firm full-time in 2005, 2006, and 2007, $N = 4,100$ firms with 90,298 attached incumbents total. Figure displays year-specific coefficients from regressions of the specified outcome on the interaction between the idiosyncratic demand shock S_j and an indicator for each year, with all interactions estimated jointly as in equation (13). Estimates for each outcome are from separate regressions. Figure plots confidence intervals based on standard errors clustered at the firm level to account for potential serial correlation of errors. Outcome is the average log monthly base wage of full-time workers, taken across all workers in the cohort of individuals who were incumbent to the affected firm in 2007. Cohort averages in “Any Firm” specification include observations of workers at different firms, if they have moved firms since 2007. In “Same Firm” specification, averages are taken over workers remaining at the firm in the observation year. When workers do not appear in the data, they are not included in the observation-year average, therefore, the underlying sample of workers is not balanced. Firms are weighted by the total number of attached incumbents present in 2007, weights are fixed across years. All regressions include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level.


Figure 8: Heterogeneity in Reduced-Form Wage Effects

Relationship Durability Defined As:



Notes: Figure shows heterogeneity in reduced-form effects of the idiosyncratic shock S_j on the wages of individuals who worked full-time at shocked firms in 2005-2007, based on two different measures of industry-level employment relationship durability. Sample is all firms in analysis sample with at least one attached incumbent, defined as workers who were present at the treated firm full-time in 2005, 2006, and 2007, $N = 4,100$ firms with 90,298 attached incumbents total. Outcome is the average log monthly base wage of full-time workers, taken across all workers in the cohort of individuals who were incumbent to the affected firm in 2007, averages include observations of workers at different firms, if they have moved firms since 2007. Figure displays year-specific coefficients from regressions of the cohort-level outcome on the interaction between the idiosyncratic demand shock S_j and an indicator for each year, separately for firms high and low relationship durability industries. All interaction coefficients for both subgroups are estimated jointly as in equation (16). The sample is split into high and low employment relationship durability groups based on whether a firm's industry's average of one of two measures is above or below the median for the sample. The two measures are either the out-of-sample sample five-digit industry average of median tenure of permanent contract workers in 2003-2007, or the out-of-sample average annual separation rate of permanent contract workers (averaged across years)—the relevant measure is specified in the table heading. All regressions include year-specific controls for the relevant "high durability" indicator. Firms are weighted by the total number of attached incumbents present in pre-period, weights are fixed across years. Standard errors are clustered at the firm level. All specifications include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years.

Figure 9: Heterogeneity in IV Pass-Through Rates



Notes: Figure shows heterogeneity in IV pass-through elasticities, which measure the change in log wages of individuals who worked full-time at shocked firms in 2005-2007 that corresponds to a one-unit change log output caused by the idiosyncratic shock S_j . Sample is all firms in analysis sample with at least one attached incumbent, defined as workers who were present at the treated firm full-time in 2005, 2006, and 2007, $N = 4,100$ firms with 90,298 attached incumbents total. Figure displays estimates of subsample-specific pass-through elasticities, obtained from the interacted difference-in-difference regression specification in (17); estimates are the same as those in Table 9. The sample is split into high and low employment relationship durability groups based on whether a firm's industry's average of one of two measures is above or below the median for the sample. The two measures are either the out-of-sample sample five-digit industry average of median tenure of permanent contract workers in 2003-2007, or the out-of-sample average annual separation rate of permanent contract workers (averaged across years)—the relevant measure is specified in the table. Interacted coefficients are estimated jointly, where interactions of the endogenous independent variable (R_j) with the heterogeneity indicator are instrumented by interactions of the same indicator with S_j . All interacted specifications include controls for the categorical indicators times $Post_t$. Output measure is either log total sales or log value added, as specified; value added is calculated as total factor payments (labor costs plus firm earnings before interest, depreciation, amortizations, and taxes). Firms are weighted by the total number of attached incumbents present in pre-period, weights are fixed across years. Standard errors are clustered at the firm level. All specifications include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years.

Tables

Table 1: Summary Statistics: Pre-period Exports of Sample Firms

	Analysis Sample					Large Exporters	
	Mean	P25	P50	P75	SD	Mean	P50
2007 Exports, Euros	1,132,340	41,718	321,375	1,187,745	2,364,478	18,500,000	5,312,642
Pre-Period Export Exposure	0.342	0.027	0.201	0.650	0.345	0.464	0.464
# Destination Countries	5.21	2	3	6	5.7	14.27	10
# Major (2D) Products	3.6	1	2	4	4.9	6.9	4
# Detailed (6D) Products	10.3	2	4	10	22.0	25.1	11
Number of Firms	4,173					934	

Notes: Table displays export statistics for firms appearing in the export data in each of 2005, 2006, and 2007, tabulated from the firm-product-destination-year level data. Analysis sample contains all such firms with 100 or fewer employees during pre-period (2005-2007 average). Large exporters are remainder of firms with over 100 employees. Exports are measured in constant 2007 Euros. Exports/Sales is the ratio of total exports to total sales from the balance sheet data (a distinct source) averaged across years 2005-2007. Counts of destination countries and products (HS2 and HS6) pool 2005, 2006, and 2007 exports of each firms, to reflect construction of the exposure weights in (9).

Table 2: Comparison of Firms and Workers in Sample and Population

	Analysis Sample					Full Count Data				
	Mean	P10	P50	P90	SD	Mean	P10	P50	P90	SD
<i>Firms</i>										
Employees	27.9	5	21	64	23.16	3.75	0	0	6	51.67
Sales / Worker if Emp>0, Euros	173,226	33,803	103,085	374,312	238,679	98,055	0	29,892	180,425	2,894,556
Value Added / Worker if Emp>0, Euros	34,769	11,935	26,575	61,969	40,939	20,397	0	11,019	42,658	203,807
N Firms, Emp>0			4,173					278,226		
N Firms			4,173					714,212		
<i>Workers:</i>										
Monthly Wage, Euros	760.80	403	575	1310	533.58	753.12	403	561.74	1,318	530.46
Hourly Wage, Euros	4.45	2.34	3.35	7.67	3.14	4.52	2.33	3.33	8.15	3.27
Log Monthly Wage	6.49	6.00	6.35	7.18	0.49	6.47	6	6.33	7.18	0.50
Log Hourly Wage	1.34	0.85	1.21	2.04	0.49	1.35	0.85	1.20	2.10	0.50
Fixed Term, Percent of Sample	0.20					0.25				
Tenure, Months (All Workers)	121	9	93	260	105.82	90.01	4.00	57	227	98.58
Female, Percent of Sample,	0.44					0.43				
Regular Hours Per Month	171	162	173	176	7.55	168	154	173	176	11
N Workers			115,526					2,490,452		

Notes: Table compares firms and workers in the analysis sample to the population data for year 2007. All figures in currency units are in 2007 Euros. Data on firm sales and output is from the balance sheet data (IES) with N=278,226. Employment counts are tabulated from the matched employer-employee data (QP) for firms that appear in both the QP and IES, N=278,226. Worker statistics are tabulations of employee records in the cleaned QP, restricting to the highest paying full-time job (more than 120 hours per month) per worker, N=2,490,452. "Monthly Wage" and "Hourly Wage" are the base contract pay earned during regular hours during the reference month in the QP, excluding overtime, fringe payments, and bonuses. "Percent FTC" is the percent of workers at the firm with fixed-term contracts.

Table 3: Effects on Firm Sales and Output

Panel A: Effect of Shock to Export Demand S_j

	Log Exports		Any Exports		Log Total Sales		Log Value Added		Any Sales	Any Employees
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Shock x Post	0.475 (0.15)**	0.555 (0.16)**	0.015 (0.036)**	0.016 (0.025)**	0.161 (0.064)**	0.234 (0.054)**	0.177 (0.069)**	0.186 (0.057)**	-0.007 (0.028)	-0.005 (0.03)
Mean Change	-0.093	0.002	-0.214	-0.08	-0.175	-0.041	-0.129	-0.028	-0.088	-0.123
Full Sample	x		x		x		x		x	
Never-Exitters		x		x		x		x		x

Panel B: Effect of Shock to Sales Demand $S_j \times \frac{Exports_j^{pre}}{Sales_j^{pre}}$

	Log Total Sales		Log Value Added	
	(1)	(2)	(3)	(4)
Shock x ExpShare x Post	.491 (0.308)**	0.944 (0.229)**	0.674 (0.282)**	0.948 (0.246)**
Full Sample	x		x	
Never-Exitters		x		x

Notes: Sample is either full analysis sample ($N = 4,173$) or sample of firms that always report positive employment (“never exitters” $N = 2,926$), as specified. Table displays regression coefficients on the interaction between the idiosyncratic shock S_j and $Post_t$, corresponding to β in equation (14) in the text. “Pre” years are 2006, 2007 (pre-period) and “Post” years 2009, 2010, 2011 (post-period). Panel A displays effects of the idiosyncratic shock S_j , which corresponds to the predicted percentage change in exports. Panel A also presents the average change in the dependent variable from pre-to-post. Panel B displays effects of S_j interacted with the pre-period share of sales in exports, corresponding to the predicted percentage change in sales due to a shock. Specifications in Panel B include controls for the post period direct effect of the export exposure $\frac{Exports_j^{pre}}{Sales_j^{pre}}$, as well as higher order terms. Firm-year observations with zeros are treated as missing when the outcome is in logs—therefore, the baseline sample is not a balanced panel, but the never-exiter sample is. Regressions are weighted by the average number of full-time employees in 2005, 2006, 2007. All regressions include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level. ** indicates $p < .05$, * indicates $p < .10$.

Table 4: Effects on Labor Adjustment

	% Change in Employment Since 2007						Average Log Monthly Wage, Current Workers	
	Log Payroll (1)	Log Employees (2)	Log Hours Per Worker (3)	Total (4)	Post-2007 Hires (5)	Incumbent Retention (6)	All Workers (7)	Post-2007 Hires (8)
<i>All Firms:</i>								
Shock x Post Effect:	0.111 (0.044)**	0.076 (0.041)*	0.011 (0.007)*	0.057 (0.041)	0.041 (0.025)*	0.016 (0.027)	0.027 (0.011)**	-0.036 (0.048)
Mean Change	0.020	-0.077	0.000	-0.139	.138	-0.276	0.105	0.455
<i>Never-Exiters:</i>								
Shock x Post Effect:	0.126 (0.043)**	0.095 (0.038)**	0.005 (0.006)	0.084 (0.032)**	0.069 (0.030)**	0.014 (0.017)	0.024 (0.013)*	-0.069 (0.053)
Mean Change	.070	-0.029	0.000	0.010	0.177	-0.170	0.103	0.453

Notes: Sample is either full analysis sample ($N = 4,173$) or sample of firms that always report positive employment ("never exiters" $N = 2,926$), as specified. Table displays regression coefficients on the interaction between the idiosyncratic shock S_i and $Post_t$, corresponding to β in equation (14) in the text. "Pre" years are 2006, 2007 (pre-period) and "Post" years 2009, 2010, 2011 (post-period). Table displays effects of the idiosyncratic shock S_i , which corresponds to the predicted percentage change in exports. Firm-year observations with zeros are treated as missing when the outcome is in logs—therefore, the baseline sample is not a balanced panel, but the never-exiter sample is. Log Hours Per Worker is average log hours (including overtime) for all workers present at the firm in the observation year. Percent changes in employment hold the denominator fixed as the 2007 employment at the firm; the hires and retention results are scaled by the same denominator. "Post-2007 Hires" in column 5 is the change in employment over 2007 accounted for by hires since 2007 that remain at the firm at the outcome year. "Incumbent Retentions" is the percentage of 2007 incumbents present at the firm in the observation year. Average Log Monthly Wage, Current Workers is the average log wage of full-time employees (regardless of contract or tenure) of workers currently at the firm; "Post-2007 Hires" in column 9 excludes workers who were present at the firm in 2007. Table also presents the average change in the dependent variable from pre-to-post. Regressions are weighted by the average number of full-time employees in 2005, 2006, 2007. All regressions include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level. ** indicates $p < .05$, * indicates $p < .10$.

Table 5: Test of Effects of Idiosyncratic and Common Components on Other Firms

	Effect on Mean Outcome Level for Other Firms in Same:					
	5-Digit Industry and Municipality		5-Digit Industry		Municipality	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Sales	0.0393 (0.0498)	1.067 (0.223)**	-0.079 (0.101)	1.610 (0.317)**	-0.012 (0.007)	0.040 (0.025)
Log Value Added	-0.022 (0.072)	1.798 (0.502)**	-0.112 (0.098)	1.981 (0.387)**	-0.007 (0.007)	0.010 (0.016)
Log Payroll	-0.009 (0.027)	0.452 (0.080)**	-0.107 (0.081)	0.706 (0.243)**	-0.008 (0.004)	0.019 (0.014)
Log Employees	-0.008 (0.022)	0.300 (0.064)**	-0.079 (0.032)	.274 (0.092)**	-0.004 (0.003)	0.010 (0.009)
Shock Component	Idiosyncratic Common		Idiosyncratic Common		Idiosyncratic Common	
N Treated Firms	3885	3885	4166	4166	4166	4166

Notes: Sample includes all firms in analysis sample with at least one other firm in the same group; the number of sample firms included in each specification is reported in the table. Table displays results from regressions where the treatment is either the idiosyncratic shock to each sample firm j , S_j , or the common component of demand, C_j , and the outcome is the leave-one-out (employment-weighted) average outcome level of all other firms (excluding j) in the same specified group as j . Entries are regression coefficients on the interaction between the idiosyncratic shock S_j and $Post_t$, corresponding to β in equation (14) in the text; each estimate is from a separate regression. Regressions are run at treated firm level; regressions are weighted by firm- j employment to match specification in Table 3. Specification includes year and treated-firm-group fixed effects; treated-firm controls are omitted, though results are similar if included. Municipalities are 24 subregions in Portugal, 5-digit industry code is the most detailed NACE classification available. Standard errors are clustered at the category-by-year level. ** indicates $p < .05$, * indicates $p < .10$.

Table 6: Pass-Through Elasticity: Effect in Wages of Attached Incumbents for Given Change in Output

	Pass-Through of Output Measure to:									
	Log Monthly Wage, Any Job		Log Hourly Wage, Any Job		Log Monthly Wage, if Same Firm		Log Hourly Wage, if Same Firm		Log (1 + Hourly Wage)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
IV Elasticity	0.137 (0.053)**	0.158 (0.068)**	0.128 (0.052)**	0.149 (0.066)**	0.131 (0.055)**	0.152 (0.070)**	0.123 (0.055)**	0.141 (0.070)	0.098 (0.042)**	0.113 (0.053)**
<i>First stage F</i>	73.79	33.56	73.79	33.56	73.55	37.76	73.66	33.56	73.79	33.56
OLS Elasticity	0.027 (0.004)**	0.02 (0.002)**	0.022 (0.003)**	0.017 (0.002)**	0.021 (0.004)**	0.017 (0.002)**	0.017 (0.003)**	0.014 (0.002)**	0.018 (0.003)**	0.014 (0.002)**
<i>Output Measure:</i>										
Log(Sales)	x		x		x		x		x	
Log(VA)		x		x		x		x		x
Mean Pre-Post Change	0.096		0.110		0.103		0.114		0.086	

Notes: Sample is all firms in analysis sample with at least one attached incumbent, defined as workers who were present at the treated firm full-time in 2005, 2006, and 2007, $N = 4,100$ firms with 90,298 attached incumbents total. Table displays estimates of pass-through elasticities, obtained from difference-in-difference regressions of average incumbent log monthly wages on the specified outcome variable; elasticity is coefficient $\epsilon^{w|R}$ on the interaction of the output measure R_j with the post period indicator $Post_t$ in (15). Table displays elasticities obtained both from OLS estimation (15) and instrumental variables (two stage least squares) estimation using the idiosyncratic shock S_j as an instrument for the output measure R_j . Output measure is either log total sales or log value added, as specified; value added is calculated as total factor payments (labor costs plus firm earnings before interest, depreciation, amortizations, and taxes). Outcomes are the average of the specified wage metric taken across all workers in the incumbent cohort; outcome averages in Columns 1-4 and 9-10 include observation-year wages at different firms, if employers have moved; outcomes in columns 5-8 are averages taken over workers remaining at the firm in the observation year. When workers do not appear in the data, they are not included in the observation-year average, therefore, the underlying sample of workers is not balanced in Columns 1-8; Columns 9 and 10 treat missing values as zeros for calculation of $\log(1+\text{wage})$. Firms are weighted by the total number of attached incumbents present in pre-period, weights are fixed across years. Baseline pre-to-post changes are displayed, all figures are nominal. All specifications include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level. ** indicates $p < .05$, * indicates $p < .10$.

Table 7: Robustness of Pass-Through Estimates

<i>IV Elasticity with Respect to:</i>	Outcome: Log Monthly Wage, Any Job									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log Sales	0.137** (0.053)	0.096** (0.044)	0.108** (0.046)	0.145** (0.062)	0.132 (0.081)	0.146 (0.09)	0.117** (0.049)	0.041 (0.158)	0.149* (0.082)	0.141 (0.088)
Log Value Added	0.158** (0.068)	0.117** (0.057)	0.151** (0.072)	0.167** (0.078)	0.159 (0.105)	0.143 (0.091)	0.140** (0.066)	0.007 (0.06)	0.150* (0.079)	0.148* (0.086)
N firms	4,100	4,100	4,100	4,100	3,993	2,578	-	5,008	4,100	5,008
N workers	90,298	90,298	90,298	90,298	87,846	67,504	-	303,095	90,298	303,095
Baseline Controls	x		x	x	x	x	x	x	x	x
Pre-Period Attribute Controls			x							
Destination Controls				x						
5 Digit Industry FE					x					
Manufacturing Only						x				
Always-Employers Only							x			
Including Large Firms								x		x
Firm-Weighted									x	x

Notes: Table displays robustness of instrumental variables estimates in Table 6 to alternative specifications. Column 1 displays IV estimates from Columns 1 and 2 of Table 6, see table notes for details. Baseline controls are year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. “Pre-period attribute controls” include controls for year-specific effects of 2005-2007 average employment, sales, assets, hiring, labor productivity, wage levels, and fixed term contract employment. Destination controls include the share of pre-period exports going to each of 10 top destination countries, as well as predicted demand using 2003-2007 changes in imports at baseline destinations. 5-digit industry FE includes industry-by-year fixed effects for 5-digit industry codes, the most detailed NACE classification available. “Manufacturing only” includes all NACE codes below 40, inclusive of agriculture and mining industries. “Firm-weighted” indicates estimation of unweighted regressions. Standard errors are clustered at the firm level. ** indicates $p < .05$, * indicates $p < .10$.

Table 8: Pass-Through Elasticity: Subgroups of Workers

<i>Subgroup:</i>	<u>Subgroup Specific Pass-Through Onto:</u>				<i>N</i> <i>workers</i>	<i>N</i> <i>firms</i>
	Log Monthly Wage, Any Job (1)	Log Monthly Wage, Any Job (2)	Log Monthly Wage, Stayers (3)	Log Monthly Wage, Stayers (4)		
Attached (At firm 2005-2007)	0.137 (0.053)**	0.158 (0.068)**	0.131 (0.055)**	0.152 (0.070)**	90,417	4,114
Permanent Contract	0.128** (0.055)	0.145** (0.069)	0.140** (0.061)	0.159** (0.076)	90,629	4,094
Fixed Term Contract	0.173 (0.111)	0.166 (0.116)	0.148 (0.118)	0.146 (0.121)	23,300	2,926
All Workers	0.228** (0.111)	0.203** (0.092)	0.138** (0.064)	0.148* (0.077)	116,368	4,173
Very Attached (Tenure > 10 Yrs)	0.146*** (0.058)	0.204** (0.097)	0.163*** (0.061)	0.228** (0.105)	44,548	3,262
<i>Permanent Contract and:</i>						
Female	0.173* (0.092)	0.174* (0.105)	0.185* (0.102)	0.191 (0.118)	98,881	3,787
Male	0.082 (0.055)	0.095 (0.067)	0.082 (0.058)	0.096 (0.071)	50,748	3,915
Wage ≤ Firm Median	0.135** (0.064)	0.140* (0.072)	0.141** (0.067)	0.142* (0.074)	46,688	3,954
Wage > Firm Median	0.153** (0.060)	0.188** (0.083)	0.164** (0.065)	0.204** (0.092)	43,734	4,114
<i>Output Measure:</i>						
Log(Sales)	x		x			
Log(VA)		x		x		

Table display pass-through elasticities corresponding to the instrumental variables estimates in Columns 1 and 2 of Table 6 pertaining to alternative types of workers besides “attached incumbents”. East estimate is obtained from a separate equation. Specification is identical to that in Columns 1 and 2 of Table 6, except outcome is average log monthly wage taken over all workers in the incumbent cohort who are members of the specified subgroup. Sample in each specification includes all firms in analysis sample with at least one 2007 incumbent full-time worker in the stated group. Outcomes are the average of the specified wage metric taken across all workers in the incumbent cohort who are in the stated group, firms are weighted by the number of 2007 incumbents in this group. Weights are fixed across years. Number of firms and incumbent workers included in each specification are displayed in Columns 5 and 6. All specifications include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level. ** indicates $p < .05$, * indicates $p < .10$.

Table 9: Heterogeneity by Industry Relationship-Durability

<i>Pass-Through to Wage if:</i>	By Separation Rate in Industry				By Typical Tenure in Industry			
	Any Job		Same Job		Any Job		Same Job	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Relationship Durability	0.323** (0.143)	0.278** (0.128)	0.316** (0.146)	0.273** (0.130)	0.166*** (0.064)	0.175** (0.074)	0.169** (0.066)	0.177** (0.076)
Low Relationship Durability	0.031 (0.049)	0.038 (0.075)	0.024 (0.053)	0.026 (0.081)	0.058 (0.089)	0.080 (0.194)	0.016 (0.093)	-0.015 (0.198)
Coeffs Equal, p-value	0.0438**	0.102	0.0508*	0.105	0.302	0.634	0.164	0.352
<i>Output Measure:</i>								
Log(Sales)	x		x		x		x	
Log(VA)		x		x		x		x

Notes: Sample is all firms in analysis sample with at least one attached incumbent, defined as workers who were present at the treated firm full-time in 2005, 2006, and 2007, $N = 4,100$ firms with 90,298 attached incumbents total. Table displays estimates of subsample-specific pass-through elasticities, obtained from the interacted difference-in-difference regression specification in (17). The sample is split into high and low employment relationship durability groups based on whether a firm's industry's average of one of two measures is above or below the median for the sample. The two measures are either the out-of-sample sample five-digit industry average of median tenure of permanent contract workers in 2003-2007, or the out-of-sample average annual separation rate of permanent contract workers (averaged across years)—the relevant measure is specified in the table. Interacted coefficients are estimated jointly, where interactions of the endogenous independent variable (R_j) with the heterogeneity indicator are instrumented by interactions of the same indicator with S_j . All interacted specifications include controls for the categorical indicators times $Post_t$. Output measure is either log total sales or log value added, as specified; value added is calculated as total factor payments (labor costs plus firm earnings before interest, depreciation, amortizations, and taxes). Firms are weighted by the total number of attached incumbents present in pre-period, weights are fixed across years. All specifications include year fixed effects, as well as controls for year-specific effects of 2005-2007 exports, log exports, the export share of sales, and the share of exports going to Spain or Angola in those years. Standard errors are clustered at the firm level. ** indicates $p < .05$, * indicates $p < .10$.

A Theory Appendix

A.1 Deriving Imperfect Substitutability from Incumbent Replacement Costs

This appendix illustrates how the production function in (1), in which incumbent labor and external labor do not enter symmetrically, can arise from a standard production function using homogenous labor with firing and hiring costs.

Consider the following version of the model above. A firm (j) that is small relative to the labor market produces revenue R_j using an mass L of workers—regardless of recency of status as a hire or incumbent—according to the revenue function $R_j = P_j \times \bar{P} \times A_j \times \tilde{f}(L)$. Output is produced using a increasing, concave function of labor $\tilde{f}_L > 0, \tilde{f}_{LL} < 0$. As before, for the following exposition, we consider behavior when A_j and $\bar{P}$ are fixed and suppress these variables. Firms inherit an stock of incumbents L^0 , as well as an initial firm-specific demand level P_j^0 . Firms can adjust their new employment level to a choice level L^* by hiring an additional mass L^{out} of workers. We again assume firms always want to make positive hires (assuming, for example, that L^0 reflects some retirement/attrition of worker who firms intend to replace in the absence of a shock).⁶²

While addition hires can be costless obtained on the external market, replacing incumbents is more disruptive then other hiring. Let L^{inc} denote the number of *retained* incumbents. Noting that the firm wants to *increase* its employment level relative to L^0 , if firms choose to only retain L^{inc} workers (or $L^0 - L^{inc}$ workers quit unexpectedly) firms must spend a share $\tau(L^0 - L^{inc})$ of their production time recruiting and training replacements. Let $R_j^* = P_j \times \tilde{f}(L^*)$ denote the production level given chosen employment L^* , the total cost is $\tau(L^0 - L^{inc})R_j^*$; ⁶³ This cost is strictly increasing everywhere in its argument $\tau' > 0$, and $\tau(0) = 0$.

These replacement frictions reflect a combination of the three kinds of barriers [Doeringer and Piore \(1971\)](#) cite as separating internal and external labor markets: costs of rapidly replacing firm-specific know-how acquired through experience [Becker \(1962\)](#); costs of having to find replacements for good incumbent matches [Jovanovic \(1979b\)](#); or direct institutional barriers to firing due to laws or unions ([Lazear, 1990](#); [Bertola and Cabellero, 1994](#)).⁶⁴ Each of these reflects an opportunity cost to replacing incumbents—forgoing production in order to retrain and recruit is costlier when product market conditions are better.

Firm profits are as follows

$$\Pi_j = (1 - \tau(L^0 - L^{inc})) \times P_j \tilde{f}(L^{inc} + L^{out}) - \bar{w}L^{out} - w_j \times L^{inc} \quad (\text{A.1})$$

⁶²While the model is not dynamic, a similar feature would occur when considering a shock to a dynamic steady state—given steady-state P_j , firms always hire $h = \delta L^0$ to maintain steady-state L . Under the assumption that shocks are never so negative so that the constraint that $h \geq 0$ binds, firms would always adjust employment by adjusting hiring levels.

⁶³This specific functional form of the cost is not important, but the implied assumption that the costs of adjustment are higher when P_j is higher is critical, as will become apparent below.

⁶⁴For example, if employees can bring suit after a dismissal, or if unionized incumbents can directly slow production at a firm via disruptive striking, then firms may lose a fraction of their output due to resulting shutdowns.

Rewriting $(1 - \tau(L^0 - L^{inc}))\tilde{f}(L^{inc} + L^{out}) = f(L^{inc}, L^{out})$, we can rewrite profits as

$$\Pi_j = P_j \times f(L^{inc}, L^{out}) - \bar{w}L^{out} - w_j \times L^{inc}$$

Thus, the model is exactly as in Section 2.⁶⁵

A.2 Proof of Proposition 1

Proposition one stated that under the assumptions of the model, when $f_1(L^{inc}, L^{out}) > f_2(L^{inc}, L^{out})$, the quasi-rent-sharing term $\rho(P_i)$ always satisfies $\rho(P_i) > 0$ and $\rho'(P_i) > 0$. We prove the result here.

Proof. The goal is to show that $V(L^{inc}, P_j) - V(0, P_j) - \bar{w}L^{inc}$ is always positive and increasing in P_j .

First, consider the derivative of the optimized quantity $V(L^{inc}, P_j) = \max_{L^{out}} \{P_j \times f(L^{inc}, L^{out}) - \bar{w}L^{out}\}$ with respect to the number of incumbents L^{inc} around some initial level of L^{inc} . Invoking the envelope theorem, $\frac{dV^1}{dL^{inc}} = P_j \times f_1(L^{inc}, L^{out*})$, where L^{out*} is the optimal choice of L^{out} . Under the assumption that $f_1(L^{inc}, L^{out}) > f_2(L^{inc}, L^{out}) \forall L^{inc}, L^{out}$, this derivative satisfies $\frac{dV^1}{dL^{inc}} = P_j \times f_1(L^{inc}, L^{out*}) > P_j \times f_2(L^{inc}, L^{out*}) = \bar{w}$, where the last equality results from the firm's first order condition for L^{out*} .

Since $\frac{dV}{dL^{inc}} > \bar{w}$ holds for all L^{inc} , reducing L^{inc} to zero from any initial level would lower V^1 by strictly more than $\bar{w} \times L^{inc}$. Thus $V(L^{inc}, P_j) - V(0, P_j) = \int_0^{L^{inc}} \frac{dV(\lambda, P_j)}{dL^{inc}} d\lambda > \int_0^{L^{inc}} \bar{w} d\lambda = \bar{w}L^{inc}$, directly implying that $\rho(P_j) > 0$ for all P_j .

Next, consider how the derivative $\frac{dV^1}{dL^{inc}}$ varies with P_j . In particular, $\frac{dV^1}{dL^{inc}dP_j} = f_1(L^{inc}, L^{out*}) > 0$ for all P_j and L^{out} . Thus, $\frac{d[V(L^{inc}, P_j) - V(0, P_j)]}{dP_j} = \frac{d}{dP_j} \int_0^{L^{inc}} \frac{dV(\lambda, P_j)}{dL^{inc}} d\lambda = \int_0^{L^{inc}} \frac{dV^1(\lambda, P_j)}{dL^{inc}dP_j} d\lambda > 0$. Since $\bar{w}L^{inc}$ is invariant to P_j , it directly follows that $\rho'(P_j) > 0$ for all P_j . $\square$

B Test of Selection of Firms into Export Markets

To motivate a test of exogenous assignment of firms to differently-shocked markets, consider a simple model of selection into export markets with better outcomes. Let Δ_{pc} be the change in total non-Portuguese imports in the market for product p in country c during the recession as defined in (10), and let the corresponding change in observed exports by the individual firm j to that same p, c market during the same period as $\tilde{E}_{j,pc}$. Suppose export growth $\tilde{E}_{j,pc}$ is determined by market level demand Δ_{pc} , unobservable firm productivity ϕ_j (denoted as such as it may reflect changes in worker productivity ϕ_i for workers i at firm j) that increases export performance of firms across

⁶⁵Technically, $f_1(L^{inc}, L^{out})$ is not defined for $L^{inc} \geq L^0$. However, since upwards adjustment of L^{inc} is not possible in the model in the main text, this restriction does not impact the analysis or the proof of Proposition 1.

all destinations in constant proportions, and an idiosyncratic residual $v_{j,pc}$ according to:

$$\tilde{E}_{j,pc} = \phi_j + \beta \Delta_{j,pc} + v_{j,pc} \quad (\text{A.2})$$


In this setting, selection occurs when firms with higher levels of ϕ_j tend to have relationships in markets with higher level of $\Delta_{j,pc}$ —that is, firms experiencing larger labor productivity shocks select into relationships with customers in markets with better import demand. When this is the case, changes in the average level $\Delta_{j,pc}$ at firm j will be correlated with unobserved heterogeneity in ϕ_j , confounding identification of idiosyncratic shocks.

Although this type of selection can not be directly tested in firm-level data, it can be partially tested by studying firms with multiple export destinations in the relationship-level data in a test similar to that used by Khwaja and Mian (2008) to test for sorting of borrower firms to lending banks that experienced differential credit supply shocks. Under the assumption of constant effects of supply productivity across destinations, no selection implies $\Delta_{j,pc} \perp \phi_j$. If this is the case, then regression estimates of β in equation (A.2) for firms with multiple pre-period destinations should be both positive (necessary condition for a causal effect of demand shocks) and invariant to inclusion of a firm fixed effect ϕ_j . If inclusion of a firm fixed effect significantly reduces the estimated magnitude of β , this would imply a positive correlation between $\Delta_{j,pc}$ and ϕ_j , contradicting the no-selection condition.


Appendix Table A.2 presents estimates of equation (A.2) with and without fixed effects. The firm-by-market sample includes all firms with at least two export destinations (across product-country cells), and only includes relationships that had positive exports in the pre-period. For the benchmark estimates, the outcome is the symmetric growth rate of exports by the firm to the specified product-by-country market from 2006-2007 to 2009-2010, which incorporates export volumes of zero in the post period (which occurs in the majority of individual relationships). Import behavior at the destination is a strong predictor of exports by the firm—though the magnitude of the estimate is small, reflecting both imprecision in the relationship-level prediction and the preponderance of zero outcomes. However, the magnitude is very similar with or without inclusion of the firm fixed effect, consistent with no sorting on latent productivity. The following columns separately display effects on the intensive margin (adjustment conditional on exports) and the extensive margin (probability of zero). The magnitude of the intensive margin (measured as the effect on the log change in exports conditional on positive flows) is substantially larger and approximately 40%, and the probability of terminating the relationship is also responsive to import demand conditions. The intensive margin effect is somewhat smaller when firm fixed effects are included, but the extensive margin response is offsettingly larger when firm fixed effects are included. These findings support the claim that demand conditions have a causal effect on firms' sales to its trading partners, and that firms do not systematically sort to destinations with better demand growth based on latent productivity trends—and, therefore, also support the use of destination-level shocks to construct an exogenous firm-level demand predictor.

Appendix Figures

Figure A.1: Comparison of Sample and Population of Firms
Panel A: Firm Size Distribution




Panel B: Industry Distribution




Notes: Figures plot distributions of firms in 2007 in the IES, “All firms” include all firms in the IES with positive employment, even if not reported in the matched employer-employee data. Employment bins are not equally sized. “Worker Distribution” plots present counts where firms are frequency weighted by the number of employees. “All” and “Sample” counts are plotted on separate axes.

Figure A.2: Recession Import Growth vs. Pre-Period Growth in Foreign Markets



Notes: Figure plots the bin-average of the y-axis variable within 20 equally-sized quantile bins of the x axis variable. Observations are product-by-country markets. Y axis variable are the arc changes in non-Portuguese imports from 2006 and 2007 to 2009 and 2010 (Δ_{pc} in Equation 10). X axis variable is 2003-2006 arc change in non-Portuguese imports in the same markets, years are chosen so the two periods do not overlap.


Figure A.3: Linearity/Symmetry of Wage Effects
Residualized Bin Scatter + Quadratic Fit



Average nominal w change in sample is 9.2%


Notes: Sample includes all firms with attached incumbent workers in 2007, $N = 4,100$. Figure plots mean residual values of the outcome within 20 equally-sized quantile bins of the residualized shock levels, where residuals are constructed to reflect all specification details of the reduced-form difference-in-difference specification in equation (14). The average slope line represents the differences-in-differences coefficient that would be obtained from the regression, with firms weighted by the number of attached incumbents in 2007. Also displayed is the best fit quadratic polynomial to the residuals given the specification details. Specification details are as in Table 7 (estimating the reduced-form effect of the shock on the outcome, rather than IV estimation using output as the explanatory variable).

Figure A.4: Firm Pay Premiums and Job Durability



Notes: Figure plots the bin-average of the y-axis variable within 20 equally-sized quantile bins of the x axis variable, including the best fit line. Sample is largest connected set of firms and full-time male workers over 18 years old who change jobs, based on years 2003-2007 in the QP. “AKM” firm effects are firm fixed effects from a wage regression with firm fixed effects, person fixed effects, year effects, and age $\times$ education group dummies. X-axis variable in “Industry Average Rate” pane is the 5-digit industry index used in Table 9 calculated for all industries in the QP, common to all firms in same 5-digit industry. X-axis variable in “Firm Rate” pane is the ratio of current-year employees not present at the firm in the following year divided by all current-year employees, averaged across 2003-2007; both X and Y variables are residuals conditional on 5-digit industry fixed effects.

Figure A.5: Effects of Idiosyncratic and Common Shocks on Exit and Job Loss



Notes: Figure shows effects of idiosyncratic shock component S_j and common shock component C_j both on firm survival and on incumbent employment and retention, whether or not firm survives. Figure displays year-specific coefficients from regressions of the specified outcome on the interaction between the idiosyncratic demand shock S_j and an indicator for each year, with all interactions estimated jointly as in equation (13). Estimates for each outcome are from separate regressions. Sample and specification for firm-level survival analysis (“any employees”) is the same as in Figure 4, except includes full sample of firms including exiters ($N = 4,173$). Sample and specifications for analysis of effects on attached incumbents are exactly as in Figure 7. “Any Employees” denotes presence of at least one full-time employee in the QP. “Any job” is a worker-level indicator denoting presence of any full-time job in the QP. “Same firm” is an indicator denoting whether employee is employed full-time at 2007 firm.

Appendix Tables

Table A.1: Comparison of Firms and Workers in Sample and Population

	Industry Heterogeneity Subsamples						Never-Exiter Firms		
	High Separation Rate			Low Separation Rate					
	Mean	P50	SD	Mean	P50	SD	Mean	P50	SD
<i>Firms</i>									
Employees	27	19	23.34	29	23	22.9	29	23	23.33
Sales/Worker if Emp>0, Euros	180,041	104,549	265,262	164,039	101,298	199,688	178,275	108,638	221,231
Value Added / Worker if Emp>0, Euros	35,250	25,224	49,015	34,030	27,524	28,558	33,765	28,409	33,764
N Firms, Emp>0		2,223			1,944			2,926	
N Firms									
<i>Workers:</i>									
Monthly Wage, Euros	747.48	545.00	540.85	774.81	600.00	524.96	769.47	582.46	541.78
Hourly Wage, Euros	4.37	3.17	3.20	4.53	3.49	3.08	4.50	3.41	3.20
Log Monthly Wage	6.46	6.30	0.50	6.51	6.40	0.48	6.40	6.27	0.44
Log Hourly Wage	1.32	1.15	0.51	1.37	1.25	0.48	1.26	1.13	0.44
Fixed Term, Percent of Sample	0.20			0.20			0.20		
Tenure, Months (All Workers)	115	90	99.75	128	95	125.83	127	99	108.07
Female, Percent of Sample,	0.51			0.36			0.43		
Regular Hours Per Month	171	173	7.45	171	173	7.66	171	173	7.56
N Workers		59,446			55,910			84,520	

Notes: Table compares firms and workers in full sample to firm subsamples examined in the main analysis. All items are exactly as in Table 2.

Table A.2: Test of Sorting for Firms with Multiple Destinations

Outcome is Change in Exports by Firm j of Product p to Country c , Measured by:

	Symmetric Growth Rate		Log Change (if >0)		Any Exports	
	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient: Change in Destination c Imports of Product p	0.0674*** (0.0150)	0.0682*** (0.0114)	0.463*** (0.0486)	0.396*** (0.0506)	0.0266* (0.0148)	0.0324*** (0.0107)
<i>Baseline Controls</i>	x		x		x	
<i>Firm Fixed Effects</i>		x		x		x
N	101,344	101,344	35,193	34,804	101,344	101,343

Notes: Tables reports results from regressions of changes in exports by firm j of product p to destination c on the change (symmetric growth rate) of imports of p to country c from all other countries during the same period, as in equation (A.2). Observations are firm-markets pairs (markets are country $\times$ 6-digit product). Sample includes all firms in primary analysis sample with exports to at least two distinct markets in the pre-period. Changes are taken from 2006-2007 to 2009-2010. When no exports occur in the post period, log values are treated as missing. Regressions are unweighted. Standard errors are two-way clustered at the firm and market level.

Table A.3: IV Pass-Through of Common Demand Shocks to Wages: Conditioning on Survival Matters

Pass-Through to Wages of Sales Changes Instrumented by:				
	Idiosyncratic Component		Common Component	
	(1)	(2)	(3)	(4)
All Firms	0.131 (0.055)**	0.123 (0.054)**	0.064 (0.030)**	0.065 (0.032)**
Never-Exiters Only	0.112 (0.050)**	0.104 (0.048)**	0.120 (0.066)*	0.144 (0.069)**
<i>Sample:</i>				
Monthly Wage	x		x	
Hourly Wage		x		x

Notes: Table compares pass-through elasticities estimated off of common component of demand C_j versus the idiosyncratic component of demand S_j , highlighting the sensitivity to conditioning on exit when studying common shocks. Table displays estimates of pass-through elasticities, obtained from difference-in-difference regressions of average incumbent log monthly wages on the specified outcome variable; elasticity is coefficient $\epsilon^{w|R}$ on the interaction of the output measure R_j with the post period indicator $Post_t$ in (15). Specifications correspond to columns (5) and (7) in Table 6, using either monthly or hourly base wages as noted; effects of idiosyncratic shock in all-firm estimates in column (1) and (2) in this table are identical to corresponding entries of Table 6. See notes to Table 6 for specification details.